GAME THEORY
A Nontechnical Introduction

by

Morton D. Davis,
City College of New York

With a Foreword by
Oskar Morgenstern

DOVER PUBLICATIONS, INC.
Mineola, New York

VISIT…

1

An Overview

The theory of games is a theory of decision making. It considers how one should make decisions and, to a lesser extent, how one does make them. You make a number of decisions every day. Some involve deep thought, while others are almost automatic. Your decisions are linked to your goals—if you know the consequences of each of your options, the solution is easy. Decide where you want to be and choose the path that takes you there. When you enter an elevator with a particular floor in mind (your goal), you push the button (one of your choices) that corresponds to your floor. Building a bridge involves more complex decisions but, to a competent engineer, is no different in principle. The engineer calculates the greatest load the bridge is expected to bear and designs a bridge to withstand it.

When chance plays a role, however, decisions are harder to make. A travel agency may want to give its customers prompt service and yet avoid excessive telephone bills. But since the agency doesn't know what future demands will be, it doesn't know how many phones to install. By using past experience and applying the laws of probability, a balance

can be struck between losses from excessive phone bills and defecting customers.

Game theory was designed as a decision-making tool to be used in more complex situations, situations in which chance and your choice are not the only factors operating. These are the situations that will concern us from now on. Let's look at a few such examples to clarify what I am saying and postpone our analysis of these kinds of problems to the later chapters.

Companies *A* and *B* intend to buy 30 and 24 typewriters, respectively. Salesman *P* represents the company that currently supplies them both; *Q* represents a competitor. Each has time to give one sales talk at either company *A* or *B*. If they visit the same company, they divide the sales at that company equally but *P* makes all the sales to the other company. If they visit different companies, each makes all the sales at the company visited. In figure 1.1 the alternatives for salesman *Q* are represented by columns and the alternatives for *P* are represented by rows. The entry corresponding to each pair of alternatives represents *P*'s *total sales* (since there are 54 typewriter sales in all, *Q*'s total sales can be deduced from *P*'s sales).

Figure 1.1

		SALESMAN *Q*	
		VISIT *A*	VISIT *B*
SALESMAN *P*	VISIT *A*	39	30
	VISIT *B*	24	42

If *Q* and *P* both visit *A*, for example, we find the entry where column "Visit *A*" meets the row "Visit *A*" is 39. *P* gets half the sales at *A* and all the sales at *B*, which is 15 + 24 = 39. It is understood that *Q* had 15 sales since the total number of sales is 54.

Generals *P* and *Q* each want control of oil deposits that *P* now controls. Thirty acres are located at *A* and 24 are located at *B*. *Q* has enough strength to invade only one location and *P* has enough strength to defend only one. Both forces are equally matched, so if *P* and *Q* go to the same place there will be a standoff; each gets half the acreage there while *P* retains control of the other location. If *P* and *Q* choose different locations, each army will control all the acreage at its own location.

Figure 1.2 indicates the total acreage that *P* will control.

Figure 1.2

		GENERAL *Q*	
		A	*B*
GENERAL *P*	*A*	39	30
	B	24	42

On the last day of a political convention, aspiring candidates *P* and *Q* will meet delegates from either states *A* or *B*. *P* is the current favorite, so if both candidates visit the same state, each will get half the delegates of that state and *P* will get all the delegates of the other one; if they visit different states, each will get the delegates of the state visited. If *A* and *B* have 30 and 24 delegates, respectively, the possible outcomes are reflected in figure 1.3.

Figure 1.3

		ASPIRING CANDIDATE *Q*	
		A	*B*
ASPIRING CANDIDATE *P*	*A*	39	30
	B	24	42

Although these three situations differ—one involves business competition, another military conflict, and the third a political campaign—they all reduce to a single problem involving the theory of games. They differ from the problems

described earlier—building a bridge and installing telephones—in one essential respect: *While decision makers are trying to manipulate their environment, their environment is trying to manipulate them.* A store owner who lowers her price to gain a larger share of the market must know that her competitors will react in kind. A thief who robs banks rather than newsstands (because that's where the money is) must be aware that the police will be asking themselves "Where would I go if I were a thief?" and act accordingly. The bridge under construction, on the other hand, has no feelings about its own safety, and the travel agency's customers were not trying to embarrass it by calling too frequently or not calling it frequently enough.

A player involved in a game with other decision-making players is in the same position as the scientist who wanted to study a monkey's behavior. After he placed the monkey in a room and gave it time to get acclimated, he looked through the observation slot—and saw the monkey's eyeball looking back at him.

In the rest of this book I will talk about certain situations that I will call *games.* In a game there will be *players* (at least two), and each will pick a *strategy* (make a decision). As a result of this joint choice—and possibly chance, a disinterested player—the result will be a reward or punishment for each player: the *payoff.* Because everyone's strategy affects the outcome, a player must worry about what everyone else does and knows that everyone else is worrying about him or her.

The words "strategy," "player," and "payoff" have roughly the same meaning here as they do in everyday language. A *player,* a participant in the game, need not be a single person. If each member of a group has exactly the same feelings about how the game should turn out, the members may be considered a single player. So a "player" may be a corporation, a county, or a football team.

A *strategy* in game theory is a complete plan of action that describes what a player will do under all possible circumstances. In ordinary usage a strategy is considered to be something clever, but nothing like that is intended here. There are poor strategies, just as there are good strategies. In the three examples depicted in figures 1.1 to 1.3, each player had two simple strategies—A and B—but in real games strategies may be so complex that they can't be written out explicitly. Also, in some real games it might seem convenient to think of a player as using several different strategies. Competing automobile companies that fix their prices every year and chess players who rethink their position after each move are two examples. But in principle you can imagine that all these decisions are merged into one to form a single strategy, and from the point of view of the theorist this is more convenient.

A complete strategy in chess, then, would start something like this: "On my first move I will move to position A. If he then moves to B, then I'll move to B'; if he then moves to C, I'll move to C'; if he . . . If after I move to A, he moves to B, I move to B'. If he then moves to Q, I'll move to Q'. . ." It is almost impossible to describe a complete strategy in detail in any real game one actually plays, and even in such a simple game as tick-tack-toe, which is played by very young children, the task is formidable. But the practical problem of writing out an entire strategy in detail shouldn't stop us from making use of the concept any more than our inability to multiply all the numbers from 1 to 1,000 shouldn't stop us from writing a formula in which this product appears.

This distinction between theory and practice is very important; the difference becomes clear if you compare a chess player's attitude toward the game of chess with that of the game theorist. Chess has been around for many centuries, and no human being or computer has come close to mastering it. In films and cartoons the bearded chess player is often

used to symbolize profound thought, and, in fact, the chess player finds the game profound and subtle. To the game theorist, however, the game is trivial. This seems an absurd position, since game theorists are not even particularly good chess players.

The apparent paradox is easily resolved, however; chess, complex as it is, is finite, so, *in principle,* every position on the board is either (a) a win for white, (b) a draw, or (c) a win for black. Given enough time, you could start from the end of the game and work forward labeling every position that can possibly arise in the game as a win, loss, or draw and finally determine whether the game itself is a win, loss, or draw. (This technique would also indicate how to enforce the win or draw.) In practice the task is hopeless, of course, so from the chess player's point of view the game is as deep as he or she believes it to be.

There is also a difference in the way we will be analyzing our "games" and the way experts analyze parlor games. If a player in a game like chess has a winning strategy, we assume he will use it; if a player is in great difficulty and needs to adopt a very subtle defense to draw, we assume she will find it. In short, we assume players always do their best.

In real life, it may make an enormous difference how you play even if you are backing a lost cause. Against an ideal player you may be defeated, but against a real person it is well known that certain strategies induce errors. Emanuel Lasker, world chess champion for many years, felt that psychology plays a very important role in the game. He often adopted a slightly inferior opening, which initially gave him a slight disadvantage, in order to disconcert his opponent. And a Russian handbook on championship chess suggests that a player should try to force an opponent into an early commitment, even if by doing so the first player obtains a slightly inferior position. In the children's game of tick-tack-toe, the

outcome will always be a draw if there is correct play on both sides, but there is a pragmatic argument for making the first move in the corner: there is only one answer to the corner move that preserves the draw, and that is a move in the center. Any other first move allows at least four adequate replies. So, in a sense, the corner move is strongest, but it is a sense that the game theorist does not recognize. Game theorists do not speak of "slightly unfavorable positions" or "commitments" or "attacks," premature or otherwise. They are incompetent to deal with the game on these terms, and these terms are superfluous to their theory. In short, game theorists, do not attempt to exploit their opponent's folly.

Since it takes no great insight to recognize the existence of folly in this world, and since the game theorist purports to be influenced by the world, why this puristic attitude? The answer is simply this: it is much easier to recognize the existence of error than to fashion a general, systematic theory that will allow you to exploit it. So the study of tricks is left to the experts in each particular game; game theorists make the pessimistic, and often imperfect, assumption that their opponents will play flawlessly.

Games like chess, checkers, tick-tack-toe, and the Japanese game of Go are called *games of perfect information* because everyone knows exactly what is going on at all times. These games offer few conceptual problems, and they won't be discussed here. In games like poker and bridge the players are, to some extent, kept in the dark, and in this sense the games are more complex. Even as trivial a game as matching pennies in which each player must choose a strategy without knowing what an adversary is doing has this added dimension of complexity.

One of the engaging properties of game theory is that many problems can be understood immediately without any special technical background. Terms such as "zero-sum" and

"prisoner's dilemma" have become part of the vocabulary of economics and the social sciences. Particularly seductive are the problems growing out of chapters 5 and 6, both for the sophisticate and beginner. Since much of the material is accessible to the reader, a number of problems have been listed at the beginning of each chapter instead of the end. Read them and think about them before your ideas are shaped by the text. They generally don't require any great quantitative skill, but they do require thought. This forethought will make the book much more challenging.

2

The Two-Person, Zero-Sum Game with Equilibrium Points

Introductory Problems

Several variations of a certain type of game are shown in figures 2.1 to 2.4. Take a moment to see what you do in each case and guess what the likely outcome will be. The material in the chapter will be much more interesting if you think about the ideas involved first.

Each of the matrices shown below represents a game. I will describe how game 1 is played in some detail; the other games are played in much the same way.

You pick a row (*A*, *B*, or *C*) and your opponent picks a column (*I*, *II*, or *III*) at the same time so neither knows when choosing what the other has picked. The number where the row and column intersect is the amount your opponent *pays*

Figure 2.1

		YOUR OUTCOME I	II	III
YOU	A	5	−2	1
	B	6	4	2
	C	0	7	−1

Figure 2.2

		YOUR OUTCOME I	II	III
YOU	A	−2	1	1
	B	−3	0	2
	C	−4	−6	4

Figure 2.3

		YOUR OPPONENT I	II	III	IV
YOU	A	−3	17	−5	21
	B	7	9	5	7
	C	3	−7	1	13
	D	1	−19	3	11

Figure 2.4

		YOUR OPPONENT I	II	III
YOU	A	2	−5	−2
	B	3	−1	−1
	C	−3	4	−4

you in dollars. So if you pick row *A* and your opponent picks column *III*, you will receive a dollar from her. (If your opponent chose *II* you would pay two dollars to her, since the number is negative.) You may assume here and throughout the book that your opponent knows the rules of the game and is as intelligent as you are. Remember that you must take your opponent's thinking into account—if you play *C* you have a chance at your greatest possible gain, 7, but will your opponent cooperate by choosing column *II*? So once again: What do you do, why, and what should the outcome of this game be?

In each of the first four games, imagine what you would do if you *knew* in advance your opponent's strategy (take each of his or her strategies in turn). If your choice depends on your opponent's choice, how do you play when you don't know what he or she will do?

Figure 2.5 depicts a similar type of game but some of the payoffs (matrix entries) have been omitted. Can you still predict what will happen even though you don't know the missing payoffs?

Figure 2.5

		YOUR OPPONENT I	II	III
YOU	A	?	?	3
	B	?	?	4
	C	7	6	5

IN FEBRUARY 1943 General George Churchill Kenney, Commander of the Allied Air Forces in the Southwest Pacific, was faced with a problem. The Japanese were about to reinforce their army in New Guinea and had a choice of two alternative routes. They could sail either north of New Britain, where the weather was rainy, or south of New Britain, where the weather was generally fair. In any case, the journey would take three days. General Kenney had to decide where to concentrate the bulk of his reconnaissance aircraft. The Japanese wanted their ships to have the least possible exposure to enemy bombers, and, of course, General Kenney wanted the reverse. The matrix entries in figure 2.6 represent the expected number of days of bombing exposure.

It is more difficult to make decisions in this kind of a game than in the games mentioned in chapter 1. The critical difference between this game and the game of chess is that here the players lack information. Both players must decide simul-

Figure 2.6

		JAPANESE CHOICE NORTHERN ROUTE	SOUTHERN ROUTE
ALLIED CHOICE	NORTHERN ROUTE	2 days	2 days
	SOUTHERN ROUTE	1 day	3 days

taneously, so neither knows the other's strategy when choosing his or her own. The analysis of this particular game is simple enough, however. At first it seems the Allies have a problem: it would be best for them to take the same route as the Japanese, but when they make their decision they don't know what that route will be. But the problem is quickly solved when you take the Japanese view. For them, the northern route minimizes their exposure *whatever* the Allies do, so their action is clear. After working this out, the Allies' decision also becomes clear: go north.

This last is an example of a two-person, zero-sum game with equilibrium points. The term "zero-sum" (or equivalently, "constant sum") means the players have diametrically opposed interests. The term comes from parlor games like poker where there is a fixed amount of money around the table. If you want to win some money, others have to lose an equivalent amount. Two nations trading make up a non-zero-sum game since both may simultaneously gain. An equilibrium point is a stable outcome of a game associated with a pair of strategies. It is considered stable because a player unilaterally picking a new strategy is hurt by the change.

A Political Example

It is an election year and the two major political parties are in the process of writing their platforms. There is a dispute between state *X* and state *Y* concerning certain water rights. Each party must decide whether it will favor *X* or favor *Y* or evade the issue. The parties, after holding their conventions privately, will announce their decisions simultaneously.

Citizens outside the two states are indifferent to the issue.

In *X* and *Y*, the voting behavior of the electorate can be predicted from past experience. The regulars will support their party in any case. Others will vote for the party supporting their state or, if both parties take the same position on the issue, will simply abstain. The leaders of both parties calculate what will happen in each circumstance and come up with the matrix shown in figure 2.7. The entries in the matrix are the percentage of votes party *A* will get if each party follows the indicated strategy. If *A* favors *X* and *B* dodges the issue, *A* will get 40 percent of the vote.

Figure 2.7

		B'S PLATFORM		
		FAVOR *X*	FAVOR *Y*	DODGE ISSUE
A'S PLATFORM	FAVOR *X*	45%	50%	40%
	FAVOR *Y*	60%	55%	50%
	DODGE ISSUE	45%	55%	40%

This is the simplest example of this type of game. Though both parties have a hand in determining how the electorate will vote, there is no point in one party trying to anticipate what the other will do. Whatever *A* does, *B* does best to dodge the issue; whatever *B* does, *A* does best to support *Y*. The predictable outcome is an even split. If, for some reason, one of the parties deviated from the indicated strategy, this should have no effect on the other party's actions. A slightly more complicated situation arises if the percentages are changed a little, as shown in figure 2.8.

B's decision is now a bit harder. If it thinks *A* will favor *Y*, it should dodge the issue; otherwise, it should favor *Y*. But the answer to the problem is in fact not far off. *A*'s decision is clear-cut and easy for *B* to read: favor *Y*. Unless party *A* is

Figure 2.8

		B'S PLATFORM		
		FAVOR *X*	FAVOR *Y*	DODGE ISSUE
A'S PLATFORM	FAVOR *X*	45%	10%	40%
	FAVOR *Y*	60%	55%	50%
	DODGE ISSUE	45%	10%	40%

foolish, *B* should realize that the chance of getting 90 percent of the vote is very slim—indeed, not a real possibility—and that it would do best to dodge the issue.

This is the same type of situation that General Kenney had to contend with. On the face of it, the northern and southern routes both seemed plausible strategies. But the rainy northern route was obviously more favorable for the Japanese, which meant that the northern route was the only reasonable strategy for the Allies.

In figure 2.9, neither player has an obviously superior strategy. In this case, both players must think a little. Each player's decision hangs on what he or she expects the other will do. If *B* dodges the issue, *A* should too. If not, *A* should favor *Y*. On the other hand, if *A* favors *Y*, *B* should favor *X*. Otherwise, *B* should favor *Y*.

Once again, the underlying structure is not hard to analyze. While at first *B* may not be clear about what should be done, it's obvious what *B* should *not* do: *B* should not dodge

Figure 2.9

		B'S PLATFORM		
		FAVOR *X*	FAVOR *Y*	DODGE ISSUE
A'S PLATFORM	FAVOR *X*	35%	10%	60%
	FAVOR *Y*	45%	55%	50%
	DODGE ISSUE	40%	10%	65%

the issue, since whatever *A* does, *B* always does better if he or she favors *X* over dodging the issue. Once this is established, it immediately follows that *A* should favor *Y* and, finally, that *B* should favor *X*. *A* will presumably wind up with 45 percent of the vote.

These two strategies—*A*'s favoring *Y* and *B*'s favoring *X*—are important enough to be given a name: *equilibrium strategies.* The outcome resulting from the use of these two strategies—the 45 percent vote for *A*—is called an *equilibrium point.*

What are equilibrium strategies and what are equilibrium points? Two strategies are said to be in equilibrium (they come in pairs, one for each player) if neither player gains by changing strategy unilaterally. The outcome (sometimes called the payoff) corresponding to this pair of strategies is defined as the equilibrium point. As the name suggests, equilibrium points are very stable. In two-person, zero-sum games, at any rate, once the players settle on an equilibrium point, they have no reason for leaving it. If, in the last example, *A* knew in advance that *B* would favor *X*, *A* would still favor *Y*; and, similarly, *B* would not change strategy if he or she knew *A* would favor *Y*. There may be more than one equilibrium point, but if there is, they will all have the same payoff.

This is true in two-person, zero-sum games, that is. Equilibrium strategies and equilibrium points can be spoken of in *n*-person and non-zero-sum games, but the arguments favoring them are considerably less appealing. In a two-person, non-zero-sum game, for example, equilibrium points need not have the same payoff; one equilibrium point may be more attractive to *both* players than another. This will be discussed in more detail in the section on the prisoner's dilemma, in chapter 5.

If they exist, equilibrium points are easy to find. In the

game shown in figure 2.9, assume that *B* knows *A*'s strategy in advance. Since *B* would choose the *minimum* value of any row *A* chooses, *A* should choose a strategy that yields the *maximum* of these *minimum* values; this value is called the *maximin,* and it is the very least that *A* can be sure of getting. In this game, if *A* plays "Favor *X*," "Favor *Y*," and "Dodge Issue," these minimum values would be 10, 45 and 10, respectively, and the maximin would be 45.

Now imagine the rules are changed so that *A* knows *B*'s strategy in advance. *A* would be expected to choose the maximum of any column, so *B* should choose the column that minimizes these maxima. That outcome is called the *minimax;* in this game the maximum values associated with the "Favor *X*," "Favor *Y*," and "Dodge Issue" strategies of *B* are 45, 55, and 65, respectively, so the minimax is 45.

If the minimax equals the maximin, the payoff is an equilibrium point and the corresponding strategies are an equilibrium strategy pair.

An equilibrium point and its corresponding equilibrium strategies are easily recognized once they are pointed out. The payoff associated with an equilibrium point is the *smallest in its row* and *the largest in its column.* In figure 2.10 (a

Figure 2.10

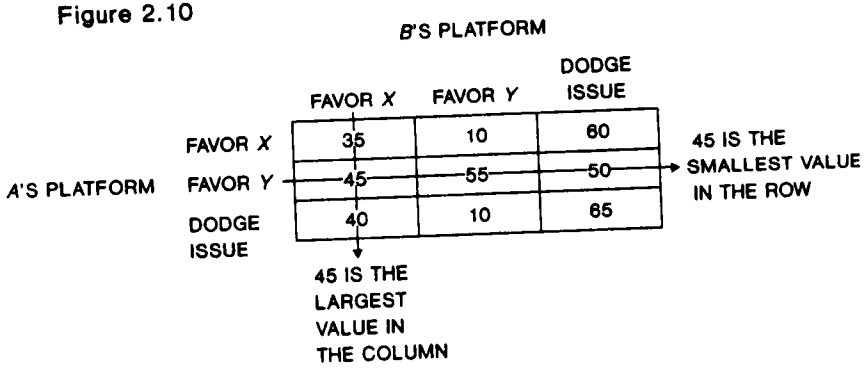

repetition of figure 2.9), the equilibrium payoff of 45 is smallest in its row and largest in its column.

When an equilibrium point exists in a two-person, zero-sum game, it is called the *solution.* Rational players should adopt the equilibrium strategies and the outcome should be the payoff associated with the equilibrium point—the *value* of the game. In the game just discussed, the equilibrium strategies for players *A* and *B* were "Favor *Y*" and "Favor *X*," respectively, and the value of the game was 45. The reasons why equilibrium points are considered solutions are:

1. *By playing his or her equilibrium strategy, a player will get at least the value of the game.* In the game shown in figure 2.10, *A* gets at least 45, whatever *B* does, if *A* plays "Favor *Y*."
2. *By playing his or her equilibrium strategy, an opponent can stop a player from getting any more than the value of the game.* By playing "Favor *X*," *B* can limit *A*'s payoff to 45 whatever *A* does.
3. *Since the game is zero-sum, a player's opponent is motivated to minimize the player's payoff.* When *A* gets 45, *B* gets 55; if *A* gets any more, it must be because *B* obtained that much less.

In games with equilibrium points, payoffs that are not associated with either equilibrium strategy have no bearing on the outcome. If in figure 2.10 the two payoffs of 10 and the payoffs of 60 and 65 were changed in *any way whatever,* the players should pick the old equilibrium strategies and the outcome should be the same.

GAME THEORY

Domination

It is often possible to simplify a game by eliminating *dominated* strategies.

Strategy *A* dominates strategy *B* if a player's payoff with strategy *A* is (a) always at least as much as that of strategy *B* (whatever other players do) and, (b) at least some of the time actually better than strategy *B*. Consider the game in figure 2.11, for example:

Figure 2.11

		YOUR OPPONENT		
		I	*II*	*III*
YOU	*A*	7	9	8
	B	9	10	12
	C	8	8	8

For you, strategy *B* dominates both strategies *A* and *C* because it always yields a higher payoff. Your opponent's strategy *I* dominates strategies *II* and *III* (recall that the matrix entries represent what your opponent pays you, so he or she wants them *small*). Although your opponent doesn't always do better with *I* than with *II* or *III*, he or she always does at least as well, and sometimes does better.

Whenever you analyze a zero-sum game, you may assume:

1. *You will never pick a dominated strategy*—why pick a dominated strategy when you can do at least as well using the strategy that dominates it?
2. *Your opponent will never pick a dominated strategy*—and this for the same reason that you won't.

In fact, if your opponent plays a dominated strategy, you can only be pleasantly surprised—you will do as well as you

would have had your opponent picked the dominating strategy, and you were already resigned to that.

If all strategies but one are dominated for each player, the equilibrium point(s) can be calculated. Consider the game shown in figure 2.12, for example:

Figure 2.12

		YOUR OPPONENT		
		I	*II*	*III*
YOU	*A*	19	0	1
	B	11	9	3
	C	23	7	−3

In this game none of your strategies are dominated initially; but since *III* dominates *I* for your opponent you can eliminate *I* from consideration. With *I* eliminated, *B* dominates *A* and *C*, and with *A* and *C* eliminated, *III* dominates *II*. The only undominated strategies, *B* and *III*, make up an equilibrium strategy pair and the value of the game is 3.

If we go back to our earlier examples, we find a number of dominated strategies. In figure 2.6 the northern route dominated the southern route for the Japanese. After eliminating the Japanese southern route, we eliminated the Allied southern route for the same reason. In figure 2.7 "Favor *Y*" dominated all of *A*'s other strategies and "Dodge Issue" eliminated all of *B*'s. In figure 2.8 "Favor *Y*" dominated everything for *A* and then "Dodge Issue" dominated everything for *B*. And finally, in figure 2.9, "Favor *X*" dominated "Dodge Issue" for *B*, "Favor *Y*" dominated everything for *A*, and "Favor *X*" dominated "Favor *Y*" for *B*.

If a game has an equilibrium point, it is easy to choose appropriate strategies and predict the outcome. But what if there is no equilibrium point? Take the simple game of matching pennies, for example as depicted in figure 2.13.

Figure 2.13

		OPPONENT: HEADS	OPPONENT: TAILS
YOU	HEADS	−1	+1
	TAILS	+1	−1

Since no strategy is dominated and there is no equilibrium point, it is hard to see how you can play such a game rationally. To try to construct a theory for such a game would appear to be a waste of time. Von Neumann and Morgenstern (1953), put the problem this way:

> Let us imagine that there exists a complete theory of the zero-sum two-person game which tells a player what to do and which is absolutely convincing. If the players knew such a theory then each player would have to assume that his strategy has been "found out" by his opponent. The opponent knows the theory, and he knows that a player would be unwise not to follow it. Thus the hypothesis of the existence of a satisfactory theory legitimatizes our investigation of the situation when a player's strategy is "found out" by his opponent. (p. 148)

The paradox is this: if we are successful in constructing a theory that indicates which strategy is best, an intelligent opponent with access to all the information available to us can use the same logic to deduce our strategy. The opponent can then second-guess us and win. So it would be fatal to use the "preferred" strategy suggested by the theory.

In fact, however, you can construct a theory that enables you to play such games intelligently. Such a theory will be the focus of the next chapter.

Solutions to Problems

1. Your opponent should see that your *B* strategy dominates your *A*—that is, you do better with *B whatever* your opponent does. Once your opponent rules out *A* as a possible strategy, he or she does best with *III*. You, in turn, do best with *B*. *III* and *B* are the suggested strategies, and you should receive a payoff of 2.

2. *I* and *A* are the suggested strategies, and you should pay 2.

3. *B* and *III* are the suggested strategies, and you should be paid 5.

4. *B* and *III* are the suggested strategies, and you should pay 1.

In each of games 1 through 4, both players will *not* lose if they announce their strategy in advance to their opponent.

5. Whatever the values of the missing payoffs, you can be sure of getting 5 by playing *C* and your opponent can be sure of losing no more than 5 by playing *III*. Since either one of you can enforce the payoff of 5, this is a plausible outcome of the game no matter what the missing payoff entries are.

3

The General, Two-Person, Zero-Sum Game

Introductory Problems

In the last chapter several games were described in which the fate of each player was, at least to some extent, in the hands of the opponent. But in each case there was an equilibrium point so that each player, by his or her own efforts, could get the value of the game, which was all that could reasonably be expected. The games in this chapter have no equilibrium points; if you are to win anything like the amount that you should, you must start second-guessing your opponent in earnest.

1. Real variations of poker are too complex to analyze so in actual games decisions are made by feel and experience. In

the simplified poker game described in figure 3.1, decide what strategy you and your opponent should choose (based on your intuition) and estimate the outcome.

You and your opponent each toss a coin with a "1" on one side and a "0" on the other. After examining your own coin, you may either *bid* or *pass*. If you pass, the player with the higher number wins $2 from the other—no money passes if there is a tie. If you bid, your opponent can *see* or *fold*. If your opponent folds, you win $11; if your opponent sees, the higher number wins $12 (again, ties are a standoff).

Figure 3.1

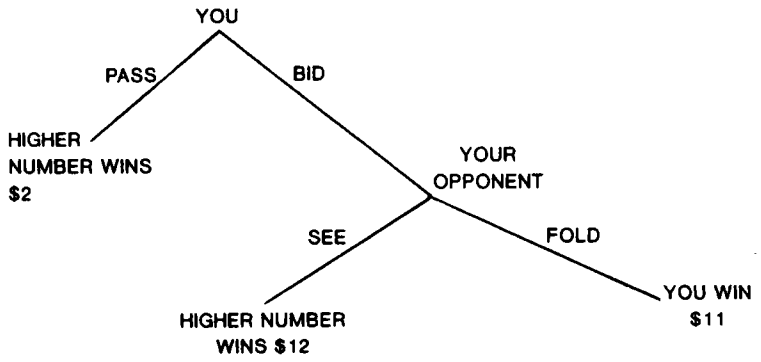

2. Two suspected smugglers are about to flee from either the sole airport or the sole seaport of a country; two police officers are assigned to catch them. If both officers guard one exit and the smugglers use the other, 100 pounds of contra-

Figure 3.2

NUMBER OF SMUGGLERS AT THE AIRPORT	NUMBER OF POLICE OFFICERS AT THE AIRPORT: 0	1	2
0	0	70	100
1	50	0	50
2	100	70	0

band will get through; if one officer guards the exit used by both smugglers, 70 pounds will get through; and if there is one smuggler at a port with no police officer, 50 pounds will get through. If there are at least as many police officers as smugglers at a port, nothing will get through. Assuming the goal of the police officers and the smugglers is to minimize and maximize the smuggled contraband, respectively, what should each do and how much will be smuggled?

3. A company's cash is contained in two safes, which are kept some distance apart. There is $90,000 in one safe and $10,000 in the other. A burglar plans to break into one safe and have an accomplice set off the alarm in the other one. The watchman has time to check only one safe; if he guards the wrong one, the company loses the contents of the other safe, and if he guards the right one, the burglar leaves empty-handed. From which safe is a sophisticated burglar more likely to steal? With what probability? What should the watchman do, and how much on average, will be stolen?
4. Your team is ahead in a team-of-four bridge tournament with one deal left. If you and your opponents in the other room reach the same contract, you win the tournament. If one team bids the unlikely slam while the other bids game, the slam bidder has a 10 percent chance of winning the tournament. What are the appropriate strategies for you and your opponents and what should the outcome be? (Note: matrix entries reflect your chance of winning the tournament.)

Figure 3.3

		YOUR OPPONENTS' BID	
		SLAM	GAME
YOUR BID	SLAM	1	.1
	GAME	.9	1

THE two-person games discussed in the last chapter were easily analyzed. They all had equilibrium points, so each player could choose a strategy that guaranteed him or her the

value of the game. In this, the chapter was misleading; all but the simplest zero-sum games that arise in practice do not have equilibrium points, and they present new and formidable problems. Edgar Allan Poe (1902) discussed such a game—an extremely simple one—in *The Purloined Letter,* and it is interesting to compare his approach with von Neumann's, which will be discussed later. Here is Poe:

"I knew one [student] about eight years of age, whose success at guessing in the game of 'even and odd' attracted universal admiration. This game is simple, and is played with marbles. One player holds in his hand a number of these toys, and demands of another whether that number is even or odd. If the guess is right, the guesser wins one; if wrong, he loses one. The boy to whom I allude won all the marbles of the school. Of course he had some principle of guessing; and this lay in mere observation and admeasurement of the astuteness of his opponents. For example, an arrant simpleton is his opponent, and, holding up his closed hand, asks, 'are they even or odd?' Our schoolboy replies, 'odd,' and loses; but upon the second trial he wins, for he then says to himself, 'the simpleton had them even upon the first trial, and his amount of cunning is just sufficient to make him have them odd upon the second; I will therefore guess odd;'—he guesses odd, and wins. Now, with a simpleton a degree above the first, he would have reasoned thus: 'This fellow finds that in the first instance I guessed odd, and, in the second, he will propose to himself, upon the first impulse, a simple variation from even to odd, as did the first simpleton; but then a second thought will suggest that this is too simple a variation, and finally he will decide upon putting it even as before. I will therefore guess even;'—he guesses even, and wins. Now this mode of reasoning in the schoolboy, whom his fellows termed 'lucky,'—what, in its last analysis, is it?"

"It is merely," I said, "an identification of the reasoner's intellect with that of his opponent."

"It is," said Dupin; "and, upon inquiring of the boy by what means he effected the *thorough* identification in which his success consisted, I received answer as follows: 'When I wish to

find out how wise, or how stupid, or how good, or how wicked is any one, or what are his thoughts at the moment, I fashion the expression of my face, as accurately as possible, in accordance with the expression of his, and then wait to see what thoughts or sentiments arise in my mind or heart, as if to match or correspond with the expression.'" (p. 123)

"As poet *and* mathematician, he would reason well; as mere mathematician, he could not have reasoned at all, and thus would have been at the mercy of the Prefect." (p. 127)

It will be easier to appreciate von Neumann's approach to this problem if we look a little more at the problem first. Put yourself in the shoes of one of the schoolboys who was pitted against Poe's remarkable wonder child.

Your situation seems hopeless; every idea that occurs to you will occur to the wonder child too. You can try to fool him by fixing your face into an "even configuration" and then playing odd. But how do you know that what you think of as an "even" configuration is not really the expression of someone playing odd and trying to look as though he were playing even? Perhaps you ought to look even, think even, and then (slyly) *choose* even. But once again, if you are clever enough to figure out this tortuous plan, won't he be clever enough to see through it? This kind of reasoning can be extended to any number of steps and the effort would be as futile as it is exhausting.

Let's look at the payoff matrix of this game in figure 3.4. The numbers in the matrix indicate the number of marbles that you are paid by the wonder child (or that you pay, if the number is negative).

Now, instead of trying to figure out just how much insight the wonder child has, suppose you assume the worst—he is so clever that he can anticipate your thinking on every count. In effect, you must announce your strategy in advance, and

the wonder child can use the information as he wishes. In this case, it seems to make little difference what you do. Whether you choose even or odd, the result will be the same: the loss of a marble. Thus, if the wonder child has perfect insight, it appears that you may as well resign yourself to the loss of a marble. You certainly can't do any worse than this—but is there a chance of doing better?

Figure 3.4

		POE'S STUDENT: EVEN	POE'S STUDENT: ODD
YOU	EVEN	−1	+1
	ODD	+1	−1

The fact is that you *can* do better, despite your opponent's cleverness. And, ironically, the means of doing so is inadvertently suggested by the second quotation from Poe: *the way to do better is not to reason at all.* To see why, let's go back a bit.

I said earlier that, as chooser, you have only these alternatives: to pick an even number of marbles or an odd number of marbles. In one sense this is true; ultimately you have to make one of these two choices. But in another sense this is *not* true. There are many *ways* in which you can make this choice. And although it may seem that *what* you decide is all that matters and *how* you decide is irrelevant, how you decide is actually critical. Of course, you can always pick one of the *pure* strategies: odd or even. But you can also use a chance device such as a die or roulette wheel to make the decision for you. Specifically, you might throw a die and if a six appears choose even and otherwise choose odd. A strategy that prescribes the selection of a pure strategy by means of a random device is called a *mixed* strategy.

From the second point of view, you have not two but an

infinite number of mixed strategies: you can choose an odd number of marbles with probability p and an even number of marbles with probability $1-p$ (p being a number between zero and 1). Roughly speaking, p is the fraction of times an odd number would be chosen if the game were played an infinite number of times. The pure strategies—"always play even," "always play odd"—are the extreme cases, where p is zero and 1, respectively.

Returning to *The Purloined Letter,* let us suppose you play even half the time; you might flip a coin and play odd whenever it comes up heads. Suppose the wonder child guesses that p is one-half, or you tell him. There is nothing more he can learn from you—you simply don't know any more. Also, unless he has powers undreamed of even by Poe, he can't possibly predict the outcome of a random toss of a coin. If you pick this mixed strategy, the outcome will be the same no matter what your opponent does; that is, each of you will win, on the average, half the time.

To summarize: you have taken a game in which you were seemingly at the mercy of a clever opponent and transformed it into one in which your opponent had essentially no control over the outcome. But if you can do as well as this when you give away your strategy, wouldn't you do better if you kept it to yourself? No, you can't do any better, and it's clear why: your opponent can do the same thing you have done. By using a random device, he too can ensure an even chance of winning. So *each* player has it within his power to see that no advantage is gained by the other.

It is conceivable that a player may do better by not "randomizing." In the game described by Poe, for example, if a player has a talent for anticipating his opponent's choices, he might try to exploit that talent. If your opponent doesn't know about randomizing (or, mistakenly, feels he is cleverer than you), it may be possible to exploit his ignorance. But this

is a doubtful advantage at best: it can always be neutralized by a knowledgeable opponent.

It should also be noted that in the odd-even game, once a player resorts to randomizing, not only can he or she not lose (on average), no matter how well an opponent plays; he or she will not gain, *however badly an opponent plays.* The more capable your opponent, then, the more attractive the randomizing procedure.

The reasoning on which the analysis of the odd-even game is based is applicable in much more complicated situations as well.

A Military Application

General X intends to attack one or both of two enemy positions, A and B; and General Y must defend them. General X has five divisions at his disposal; General Y has three. Each general must divide his entire force into two parts and commit them to these two positions, without knowing what the enemy will do. The outcome is determined once the strategies are chosen. The general who commits the most divisions to any one position will win there; if both generals send the same number of divisions, each will get credit for half a victory. Assuming that each general wishes to maximize his expected number of victories (that is, the average), what should he do? What is the expected outcome?

The entries in the matrix in figure 3.5 represent the average number of victories won by General Y. If General Y sends two divisions to A, for example, and General X does the same, General Y will gain half a victory at A and no victory at B, where he is outnumbered three to one.

Let's start by looking at the problem from General Y's

point of view. Assume he's a pessimist and he believes General X is capable of second-guessing him. It is clear that whatever strategy he picks, he will have no victories; General X will assign one more division to A than General Y did and have one extra division at B as well; that is, the maximin is 0.

Figure 3.5

GENERAL *Y*'S STRATEGIES: DIVISIONS SENT TO *A*	GENERAL *X*'S STRATEGIES: DIVISIONS SENT TO *A* 0	1	2	3	4	5
0	½	0	½	1	1	1
1	1	½	0	½	1	1
2	1	1	½	0	½	1
3	1	1	1	½	0	½

Now suppose General Y decides to play a mixed strategy; that is, he makes his decision on the basis of a chance device. We will still assume that General X can guess the nature of the chance device (the probabilities of selecting each of the strategies), but not the actual outcome of the toss of the coin or spin of the wheel. Specifically, suppose General Y assigns all his divisions to A one-third of the time, all his divisions to B one-third of the time, and makes each of the two possible two/one splits one-sixth of the time.* General X, aware of General Y's strategy, need examine only what will happen in each of his own six strategies and then select the one that gives him the greatest advantage. It is a matter of simple calculation to determine that General X will obtain an average of 1 5/12 victories (and, correspondingly, General Y will obtain an average of 7/12 victories) if he does not put all his

*I am primarily concerned with describing the meaning and significance of the von Neumann solution. (The technique of computing the strategies and value of the game will be discussed later.) For now, I will pull solutions out of a hat without apology whenever appropriate.

divisions in one location, and he will obtain an average of 1 1/6 victories otherwise.

What can we conclude from all this? Simply that General Y can reasonably expect to win an average of 7/12 victories, no matter how clever General X is. In fact, he can win the 7/12 victories even if he tells General X his strategy in advance. But should he settle for this? To answer this question, let us look at the problem from General X's point of view.

It is clear from the start that if General X commits himself to a pure strategy and General Y figures it out, they will each get one victory. When five divisions are distributed between two locations, there must be two or less divisions at one of the locations. If General Y sends all three of his divisions to the location where General X sent his smaller force, he will gain precisely one victory (and there is no way he can do any better); that is, the minimax is one.

But now suppose General X sends one, two, three, four divisions to A, with probabilities one-third, one-sixth, one-sixth, one-third, respectively, and the balance of his forces to B. (A little thought should make it obvious that it never pays to send all five divisions to one location.) It makes no difference what General Y does. The result will be the same: General Y will obtain an average of 7/12 victories.

In short, the situation is this. Either General X or General Y can, without the aid of, and without errors from, his opponent, ensure an outcome at least as favorable to himself as the outcome just mentioned: 7/12 of a victory (on average) for General Y, 1 5/12 victories for General X. The answer to the earlier question is that neither general can hope to do better than this against an informed opponent.

A Marketing Example

Two firms are about to market competitive products. Firm *C* has enough money in its advertising budget to buy two blocks of television time, an hour long each, while firm *D* has enough money for three such one-hour blocks. The day is divided into three basic periods: morning, afternoon, and evening, which I will indicate by *M, A,* and *E,* respectively. Purchases of time must be made in advance and are kept confidential. Fifty percent of the television audience watch in the evening; 30 percent watch in the afternoon; and the remaining 20 percent watch in the morning. (For the sake of simplicity, it is assumed that no person watches during more than one period.)

If a firm buys more time during any period than its competitor, it captures the *entire* audience during that period. If both firms buy the same number of hours during any one period—and this is so if neither firm buys time at all—each gets half the audience. Each member of the audience buys the product of just one of the firms. How should the firms allocate their time? What part of the market should each firm expect to get?

The payoff matrix for this game is shown in figure 3.6. The strategies for firm *C* (there are six) are indicated by two letters. *MA* means one hour of advertising in the morning and another hour in the afternoon. Similarly, the ten strategies for firm *D* are expressed by three letters. The entries in the matrix represent the percentage of the market that firm *D* captures; firm *C* gets the rest.

To illustrate how the percentage entries were obtained, assume firm *D* chooses *EMM*—one evening and two morning hours—and firm *C* chooses *EM*—one evening and one morning hour. Since firm *D* bought two morning hours to firm *C*'s one, firm *D* will capture all of the morning market of

Figure 3.6

D'S STRATEGIES	*C*'S STRATEGIES: *EE*	*EA*	*EM*	*AA*	*AM*	*MM*
EEE	75	60	65	60	50	65
EEA	65	75	80	60	65	80
EEM	60	70	75	70	60	65
EAA	40	65	55	75	80	80
EAM	50	60	65	70	75	80
EMM	35	45	60	70	70	75
AAA	40	40	30	65	55	55
AAM	50	50	40	60	65	55
AMM	50	35	50	45	60	65
MMM	35	20	35	45	45	60

20 percent. Since each firm has one evening hour and no afternoon hour, these markets of 50 percent and 30 percent, respectively, will be divided equally between the two firms. Thus firm *D* will get 60 percent of the total market if the indicated strategies are used.

Solution to the Marketing Example

Notice that *EEM* dominates the three strategies *AMM, MMM,* and *AAA,* while *EAM* dominates *EMM* and *AAM.* This means that firm *D*'s last five strategies can be effectively discarded. Once this is done, *AM* dominates *MM.* (If firm *D* plays *AAM,* firm *C* would do better with *MM* than with *AM;* but we have already decided that firm *D* won't play *AAM.*) The game, then, is reduced to one in which each player has five strategies.

One solution to this problem—and, again, I won't say how

I got it—is that firm D should play each of the strategies *EEE, EEA,* and *EAM* one-third of the time, and firm C should play *EE* six-fifteenths of the time, *AA* five-fifteenths of the time, and *AM* four-fifteenths of the time. If firm D uses its recommended strategies, it can be sure of winning on average at least 63 1/3 percent of the audience; and if firm C uses its recommended strategy, firm D won't win any more than this.

Simplified Poker

Figure 3.7

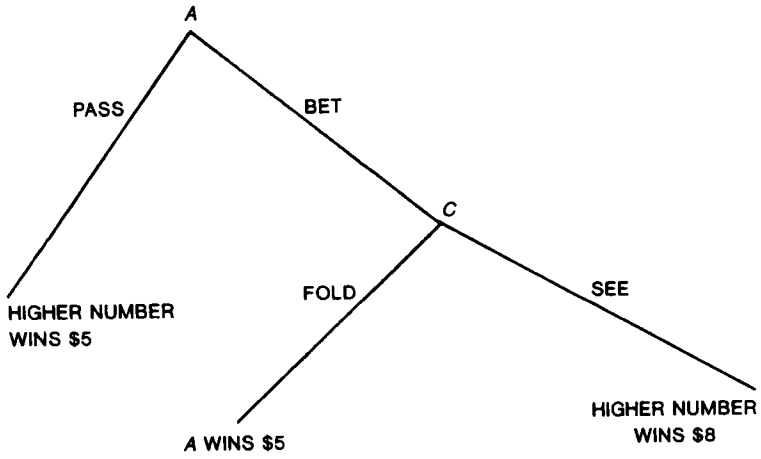

A and C each puts $5 on the table and then toss a coin which has "1" on one side and "2" on the other. Neither player knows the outcome of the other's toss.

A plays first. She may either pass or bet an additional $3. If she passes, the numbers tossed by the two players are compared. The larger number takes the $10 on the table; if both numbers are the same, each gets his or her $5 back.

If A bets $3, C can either see or fold. If C folds, A wins the $10 on the table, irrespective of the numbers tossed. If C sees, he adds $3 to the $13 already on the table. Again the numbers are compared; the larger number takes the $16, and if the numbers are equal, each gets his or her money back. What are the best strategies, and what should happen?

Each player has four strategies. A can always pass (indicated as *PP*), pass with 1 and bet with 2 (*PB*), pass with 2 and bet with 1 (*BP*), or always bet (*BB*). C can always fold (*FF*), always see (*SS*), see with 1 and fold with 2 (*SF*), and fold with 1 and see with 2 (*FS*).

To illustrate how the entries in the payoff matrix in figure 3.8 were obtained, suppose A plays *BB* and C plays *SF.* Half the time, C will get 2 and fold (A always bets), and A will win $5. A fourth of the time, C will get 1 and A will get 1 and nothing will happen. And a fourth of the time, C will get 1 and A will get 2; A will bet (as always), C will see, and A will win $8. A's average gain, if these two strategies are adopted is $4.50.

Figure 3.8

A'S STRATEGIES	C'S STRATEGIES: FF	SS	SF	FS
PP	0	0	0	0
BB	5	0	4½	½
PB	¾	¾	2	0
BP	3¾	−¾	¾	½

It can be noticed first that *BB* dominates *BP* and *PP,* while *FS* dominates *SF* and *FF.* The formal model reflects what is probably intuitively clear: whenever A gets 2 on the toss, she should bet; and whenever C gets 2 on the toss, he should see. The strategy recommended for A is to play *BB* three-fifths of the time and *PB* two-fifths of the time. This will assure A an

average win of 30 cents a play. If *C* plays *FS* three-fifths of the time and *SS* two-fifths of the time, *C* will have a loss of no more than 30 cents, on average, per play.

The Minimax Theorem

Up to now I have considered specific games only. In each of them I recommended certain strategies for the players and indicated what the outcome should be. All these examples were special instances of one of the most important and fundamental theorems of game theory: von Neumann's minimax theorem.

The minimax theorem states that one can assign to every finite, two-person, zero-sum game a value *V*: the average amount that player *I* can expect to win from player *II* if both players act sensibly. Von Neumann felt that this predicted outcome is plausible, for three reasons:

1. There is a strategy for player *I* that will protect this return; against this strategy, nothing that player *II* can do will prevent player *I* from getting an average win of *V*. Therefore, *player* I *will not settle for anything less than* V.
2. There is a strategy for player *II* that will guarantee that he or she will lose no more than an average value of *V*; that is, *player* I *can be prevented from getting away more than* V.
3. By assumption, the game is zero-sum. What player *I* gains, player *II* must lose. Since player *II* wishes to minimize his or her losses, *Player* II *is motivated to limit player* I's *average return to* V.*

The last assumption, which is easily overlooked, is crucial. In non-zero-sum games, where it does not hold, you should

*Notice that these three characteristics are almost identical to the ones possessed by games with equilibrium points that were listed in chapter 2.

not conclude that just because player *II* has the power to limit player *I*'s gains, player *I* will necessarily do so. But here, where hurting player *I* is, for *II*, equivalent to furthering his or her own ends, the assumption is compelling.

The concept of mixed strategies and the minimax theorem considerably simplify the study of these games. To appreciate how much, imagine what the theory would be like in their absence. Except for the simplest types of games—games in which there are equilibrium points—there is anarchy. It is impossible to choose a sensible course of action or predict what will happen; the outcome is intimately related to *both* players' behavior and each is at the mercy of the other's caprice. You can only try to define rational play along the lines suggested by Poe, and such attempts are virtually worthless.

With the addition of the minimax theorem, the picture is radically different. In effect, you can treat all two-person, zero-sum games as though they had equilibrium points. The game has a clear value, and either player can enforce this value by selecting the appropriate strategy. The only difference between games with actual equilibrium points and those without them is that in one case you can use a pure strategy and obtain the value of the game with certainty, while in the other case you must use a mixed strategy and you obtain the value of the game on average.

The virtue of the minimax strategy is security. Without it, you must resort to the double- and triple-cross world of Poe's precocious student. With it, you can obtain your full value, and you have the assurance that you couldn't do better—at least, not against good play.

Calculating Mixed Strategies

It is considerably harder to calculate mixed-strategy solutions than pure ones. In principle you can find solutions to any zero-sum game, but it may be difficult in practice. The following simple way of finding mixed-strategy solutions *sometimes* works.

To solve a game, first delete all dominated strategies. A strategy may be dominated by either a pure strategy or a mixed one. Domination by a mixed strategy is harder to detect, but a careful inspection of the payoff matrix and a little intuition will usually enable you to eliminate superfluous pure strategies.

In the game shown in figure 3.9, no pure strategy dominates another but your strategy *B* is dominated by the mixed strategy "play *A* one-fifth of the time and *C*, four-fifths of the time" (it is dominated by other mixed strategies as well). If your opponent plays *D*, you get $1/5(15) + 4/5(5) = 7$, and if your opponent plays *E* you get $1/5(10) + 4/5(20) = 18$. In both cases, the result exceeds what you would get using *B*.

Figure 3.9

		YOUR OPPONENT *D*	YOUR OPPONENT *E*
YOU	*A*	15	10
	B	6	15
	C	5	20

Once you eliminate all dominated strategies, you can often find the mixed-strategy solution by following this rule: *Choose a mixed strategy that will give you the same average payoff whatever your opponent does.*

Let's assume that in this game you play *A*, with probability *p*, and *C*, with probability $1-p$, and your opponent plays *D*,

with probability *r*, and *E*, with probability $1-r$. If your opponent plays *D*, your average payoff is $(15)(p) + (5)(1-p) = 5 + 10p$; it is $(10)(p) + (20)(1-p) = 20 - 10p$ if your opponent plays *E*. If these are the same, $5 + 10p = 20 - 10p$ and p equals 3/4, the payoff is 12 1/2 on average.

By a similar calculation, you find that your opponent should play each of his or her strategies half the time and the payoff is the same: 12 1/2.

Let's consider another example: A convict has two possible escape routes—the main highway and a forest—and the pursuing sheriff can cover only one of them. If the sheriff and convict take different routes, the escape is certain; if both take the highway route, the convict will certainly be caught. If both go through the forest, the chance of escape is $1 - 1/n$. The matrix entries in figure 3.10 indicate the probability of escaping.

Figure 3.10

		SHERIFF: HIGHWAY	SHERIFF: FOREST
CONVICT	HIGHWAY	0	1
	FOREST	1	$1-1/n$

Using the same calculations as earlier, we find both sheriff and convict should take the highway with probability $1/(n + 1)$. If either player follows this strategy, the convict will escape with probability $n/(n + 1)$. As n increases and with it the escape probability, the less attractive the forest route becomes to the sheriff; nevertheless that is the route the sheriff must take with increasing probability.

The same method may be applied to larger matrices; in the game depicted on figure 3.11, assume you play strategies *A*, *B*, and *C*, with respective probabilities *p*, *q*, and $1 - p - q$ and your opponent plays *D*, *E*, and *F* with respective proba-

bilities, r, s, and $1 - r - s$. Applying the rule, you would conclude that $8p + 6q + 18(1 - p - q) = 6p + 15q + 6(1 - p - q) = 14p + 3q + 6(1 - p - q)$. By playing A, B, and C, with respective probabilities of one-half, one-third, and one-sixth, you would receive an average payoff of 9, whatever your opponent did.

Figure 3.11

		YOUR OPPONENT		
		D	*E*	*F*
YOU	*A*	8	6	14
	B	6	15	3
	C	18	6	6

Your opponent would obtain a similar pair of equations—$8r + 6s + 14(1 - r - s) = 6r + 15s + 3(1 - r - s) = 18r + 6s + 6(1 - r - s)$—and conclude that he or she should play D, E, and F with respective probabilities four-sixteenths, seven-sixteenths, and five-sixteenths, with the same average payoff of 9.

Solving Two-Person, Zero-Sum Games: A Summary

To solve a two-person, zero-sum game, follow this procedure:

1. Calculate the maximin and minimax—if they are equal, you have found the proper strategies and the value of the game. If they are unequal, go on to step 2.
2. Eliminate all strategies that are dominated.
3. Assign probabilities to each of your strategies so that the outcome of the game will be the same on average whatever your opponent does. Assume your opponent does the same;

if the outcome when you use this mixed strategy is the same as the outcome when your opponent uses his or her mixed strategy and if all probabilities are nonnegative, you have solved the game.

If there is a gap between the outcomes or if some of the probabilities are negative, reexamine the game for dominated strategies; if there are none, then this method failed.

Some Second Thoughts and Applications

Most knowledgeable game theorists, when asked to select the most important single contribution to game theory, would probably choose the minimax theorem. The arguments in favor of minimax strategies are very persuasive; still, they can be overstated. And even a sophisticated writer, in a careless moment, can fall into a trap. In his book *Fights, Games and Debates* (1960), Anatol Rapoport discusses the game shown in figure 3.12. The numbers in the payoff matrix indicate what player *II* pays player *I*. The units are unimportant as long as we assume that both players would like to receive as much as they can.

Figure 3.12

		PLAYER *II*	
		A	*B*
PLAYER *I*	*a*	−1	5
	b	3	−5

Rapoport starts by making the routine calculations. He computes player *I*'s minimax strategy (play strategy *a* four-sevenths of the time and strategy *b* three-sevenths of the time), player II's minimax strategy (play strategy *A* five-sevenths of the time and strategy *B* two-sevenths of the

time), and the value of the game (an average win of five-sevenths for player *I*). He then goes on to say: "To try to evade by using a different strategy mixture can only be of disadvantage to the evading side. There is no point at all in trying to delude the enemy on *that* score. Such attempts can only backfire" (p. 160). But Rapoport's conclusion that any deviation from the minimax "can only be of disadvantage" leads to a curious paradox. To see why, observe two basic facts.

Note first that if *either* player selects the minimax strategy, the outcome will be the same: an average gain of five-sevenths for player *I*. There is only one way to obtain another outcome: *both* players must deviate from their minimax strategies.

The second point is that the game is zero-sum. A player cannot win unless he or she wins from the opponent; a player cannot lose unless there is a corresponding gain for the opponent.

Putting these two facts together, we have the contradiction. If a player deviates from the minimax, it cannot possibly be "of disadvantage" unless his or her opponent deviates as well. And, by the same argument, the opponent's deviation must also be of disadvantage. But the game is zero-sum: both players cannot lose simultaneously. Obviously, something is wrong.

Here is the answer. There is a strategy for player *I* that guarantees him or her five-sevenths, and the player can be stopped from getting any more; moreover, player *II* is motivated to stop player *I* from getting more. If player *I* chooses another strategy, he or she is gambling. If player *II* also gambles, there is no telling what will happen. The minimax strategies are attractive in that they offer security; but the appeal of security is a matter of personal taste.

As a consequence of the minimax theorem, the general,

zero-sum, two-person game has a good theoretical foundation. But, like the game of perfect information, it rarely exists in practice. The difficulty is the requirement that the game be zero-sum.

The essential assumption upon which the theory is based is the opposition of the two players' interests. To the extent that the assumption is not valid, the theory is irrelevant and misleading. Often the assumption seems to be satisfied but in reality is not. In a price war, for example, it may be to the advantage of *both* parties that prices be maintained. In a friendly game of poker, *both* players may prefer that neither suffers an excessive loss. And when bargaining, buyer and seller may have divergent interests with respect to the price, but *both* may prefer to reach some agreement. Still, some situations that are best reflected by a zero-sum model—the political arena provides us with a number of them.

Every four years there is a presidential election in the United States. What appears to be one contest is really fifty; each state has its own electoral votes that are won or lost independently from every other state, and the number of electoral votes vary from state to state. The resources with which each party conducts its campaign are money, speeches by notables, media advertising, direct mail, and so forth. These resources are limited for both parties and must be used prudently. The interests of the parties are at opposite poles—there is one president, and if he is to be chosen from one party, he clearly can't be from the other.

Steven Brams and Morton Davis (1974) apply game theory to the problem of allocating resources so as to maximize the number of electoral votes. It is not surprising that candidates tend to spend more of their resources in large states than in small ones, but it seems that they should do so even out of proportion to their size. It is suggested that states should allocate their resources in proportion to the 3/2 power of the

number of electoral votes; a state with four times as many electoral votes as another should get eight times as much resources. An attempt was made to confirm these theoretical results empirically; using campaign appearances as a criterion, it turned out that the "3/2 rule" for allocating resources was a better model for the actual allocation than a proportional model. In a different but similar problem, the same authors consider a different allocation problem: What is the proper way to allocate resources in a sequence of primaries?

The essential difference between this problem and the last one is the time element: a candidate who does well in the early primaries gains political and financial support for the later ones. In this model there is a snowball effect; money spent on the first state in the primary indirectly affects the outcome in all the others. The result is that both in theory and in practice the early primaries are better financed than the later ones. According to the theory, if every state had the same number of electoral votes and there were fifty primaries, a candidate should spend fifty times more on the first primary than it spends on the last one.

The same two authors (1978) analyze another encounter that is best modeled by a zero-sum game: that of two lawyers in a courtroom. It is assumed that there is a population of jurors in which each member has a certain propensity to acquit or convict and that the attorneys can judge this from the potential juror's background. (This is certainly a simplification of the actual situation, but it is generally accepted that there is a correlation between a person's background and his or her actions as a juror.) Both the prosecutor and the defense attorney are given a certain number of peremptory challenges—dismissals of potential jurors without cause—and must decide when to exercise them. The decisions will be based on the attractiveness of the juror being considered and the number of peremptory challenges remaining to both parties. The obvious danger in challenging what appears to be an unfavorable juror is that the next prospect may be worse and the attorney will be helpless.

R. Avenhaus and H. Frick (1976, 1977) describe a zero-sum "game" in a factory. The factory processes material that is either very expensive, such as gold or platinum, or very valuable, such as nuclear fuel. The players are an inspector and a potential diverter. The inspector keeps an inventory of the amount of material that should be there and compares it to the amount of material that is actually there. Although measurements are accurate, they are imperfect; discrepancies are inevitable even when no material is diverted. Certain costs prohibit the inspector from taking inventory too often or setting off too many false alarms, and there are countervailing costs if the inspector fails to discover a diversion or if there is an excessive delay in the discovery. The inspector must observe the discrepancy between book inventory and actual inventory and then balance these costs. The diverter clearly gains from a successful diversion but pays a heavy price if caught. While this is clearly not a strict zero-sum game—the cost to the inspector of excessive inspections is not balanced by a corresponding gain by the potential diverter—it is a close enough approximation to be useful.

In a much less serious vein Robert Bartoszynski and Madan L. Puri (1981) apply a game-theoretical model to a real game: tennis. In tennis, when a player serves, the ball must fall inside certain boundaries; if it fails to do so twice in succession, the player loses the point. If the player succeeds after failing once, there is no penalty for the failure. Most players use a hard first serve, which is difficult to return but somewhat inaccurate, and reserve a more reliable, but less effective, serve if the first serve fails. Bartoszynski and Puri investigated under what conditions a player should use the hard serve both times, neither time, or just the first time.

The authors also address the following problem: suppose you have been preparing a special serve to use against your usual tennis partner. The service is to be a surprise that will only be effective once—a slice that bounces deceptively, perhaps. At what point in the game should you use it? Bartoszynski and Puri's somewhat surprising conclusion is that it makes no difference when you use it—the effect will be exactly the same. The important thing is that you use it at some time. If you reserve its use for a situation that may not arise—only when you are leading 40 to 15, for example—you lose part of your advantage.

Another, surprising application to game theory was made some time ago by mathematician Edwin O. Thorpe and W. E. Waldman (1973). People have been gambling for a very long time, and it is very likely that, for as long a time, there have been players with systems. Some systems are based on lucky numbers, omens, using past history to predict the future, changing the amounts bet, and so forth. But since the house fixed the odds so that each individual bet was unfavorable, no sorcery could alter the inevitable loss to the long-time player. So systems players patronized the gambling houses and the gambling houses welcomed and thrived on them. Until the appearance of Edwin Thorpe. Thorpe realized that making a profit on a sequence of bets when you lose on average on each one is like losing money on each sale and trying to make it up in volume. He also realized that in a game like blackjack, the situation isn't really repeated unless the cards are dealt thoroughly after each deal. Thorpe, with the aid of a computer, devised a number of techniques for determining whether the cards played in the past improved a player's chances. Thorpe's methods were described in a best-selling book, *Beat the Dealer,* (1966) and gambling casinos tacitly acknowledged his success: they changed some of the rules, combined decks to make counting harder, and attempted to bar card counters from the tables. In a more recent paper written jointly with William E. Walden, Thorpe (1973) analyzes such games as baccarat, poker, and trente-et-quarante in which cards are exposed without being redealt and in which this information can be put to a player's advantage.

Another application of game theory to real games was made by Nesmith C. Ankeny (1981). Poker has long been considered a fertile area for applying game theory, and a number of papers have been written on the subject, including a substantial part of *Theory of Games and Economic Behavior* by von Neumann and Morgenstern (1953). But generally, only simplified forms of poker were analyzed. Ankeny, a mathematics professor at M.I.T., addresses the game of poker as it is actually played in his book *Poker Strategy: Winning with Game Theory.* He suggests a "proper balance of deception and strength" to maximize winnings. Certain aspects of bridge, such as signaling, have also been analyzed, and mixed strategies have been applied to bridge play; but no overall model, encompassing both bidding and play, has been devised.

Military tactics, an area in which the parties are involved in almost pure conflict, is another source of two-person, zero-sum games. Evasion and pursuit games and fighter-bomber duels are two applications. In addition, there is a general class of games called Colonel Blotto games, in which the players are required to allocate certain resources at their disposal. We have seen one example of this in a military context—Generals X and Y allocating their divisions—and another in a marketing context. The same problem arises when allocating salespeople to various territories or police officers to high-crime areas. An interesting application of Colonel Blotto games arose in the process of creating a model for inspection in a possible disarmament agreement. One country has fac-

tories that produce certain types of weapons—some more potent than others, some easier to hide than others, and so on. The inspecting country has to decide how to allocate its inspectors to verify that the inspected country is not exceeding its quota of weapons.

An interesting application of minimax strategies was observed by Sidney Moglower (1962). In the process of selecting crops, the farmer can be regarded as one player and the "hypothetical combination of *all* the forces that determine market prices for agricultural products" as the other player. Moglewer points out that though it is difficult to justify the farmer's implicit assumption that the universe is concerned with him or her as an individual, the farmer acts as though this were the case.

Some Experimental Studies

Although I speak of game *theory,* it is important to look at games in practice. In a speech to future jurists at the Albany Law School, Justice Benjamin Cardozo remarked, "You will study the life of mankind, for this is the life you must order, and to order with wisdom, know." The comment applies to game theory as well. Game theory has its roots in human behavior; if the theory is not related to human behavior in some way, it will be sterile and meaningless except as pure mathematics. Aside from other considerations, experiments in game theory are interesting in themselves, perhaps because people enjoy reading about what other people do. This is sufficient reason for discussing them.

Richard Brayer (1964) had subjects repeatedly play the game shown in figure 3.13. The entries in the payoff matrix indicate what the subject received from the experimenter.

The players either were told they were playing against an experienced player or they were told that their opponent would play randomly, but in fact they invariably played against the experimenter.

Figure 3.13

		EXPERIMENTER		
		A	*B*	*C*
SUBJECT	*a*	11	−7	8
	b	1	1	2
	c	−10	−7	21

It takes only a moment for subjects to deduce that they should play strategy *b* if they feel their opponent is intelligent. From the point of view of the experimenter, strategy *C* is dominated by *B,* so subjects can assume that the experimenter will never play *C.* Having decided this, subjects should never play *c,* since they always do better with strategy *b.* Eliminating *c* and *C,* the experimenter should play *B,* and subjects should play *b.* (*B, b*) are equilibrium strategies, and the payoff 1 is an equilibrium point.

What actually happened? For one thing, the subjects ignored what they were told about their opponent and responded only to how the opponent played. Whether they did not see the relevance of what they were told about their opponent, or did not know how to apply it, or whether they were reacting with the natural skepticism subjects have toward anything they are told by an experimenter is not clear. The subjects did play *b* if the experimenter played *B,* but not otherwise. Postexperimental interviews confirmed what the play indicated: the players couldn't anticipate the experimenter's choice of strategy *B.* In fact, *more than half the subjects felt that the experimenter was stupid for playing* B *and settling for a loss of 1.* When the experimenter picked

a strategy at random, the subjects generally responded by playing strategy *a:* the one that gave them the highest average return.

The same pattern was observed by other experimenters. Several—Theodore Caplow (1956, 1959), Merrill M. Flood (1952), Oliver L. Lacey and James L. Pate (1960), Bernhardt Lieberman (1960), Robert E. Morin (1960), and Anatol Rapoport and Carol Orwant (1962)—concluded that most subjects are simply unable to put themselves in the shoes of their opponents. Subjects tend to prefer strategies that yield an apparently high average return (against an opponent playing at random) to equilibrium strategies. When players are punished for deviating from the equilibrium strategy, they change their behavior, but not otherwise.

In games without equilibrium points, the players have even less insight. In an interesting review of experimental games, Anatol Rapoport and Carol Orwant (1962) suggest that equilibrium points, when they exist, will be found eventually, if not immediately, by the players. The speed with which they are discerned varies with the experience and sophistication of the players and the complexity of the game. When there are no equilibrium points, when what is called for is a mixed strategy, the situation is much worse. Not only are all but the most sophisticated players incapable of making the necessary calculations, but almost all players do not even see the need for them.

Is the failure to play an equilibrium strategy, where one exists, necessarily proof of a player's ignorance? Perhaps not on the face of it. It might be argued that the case for the equilibrium strategy is based to some extent on the assumption that one's opponent will play reasonably well; if this is not true, a player often does better by not playing the equilibrium strategy. When the experimenter played the equilibrium strategy, the subjects did too. When the experimenter

played randomly and "irrationally," however, the subjects persisted in playing nonequilibrium strategies.

All this sounds plausible but, as mentioned earlier, it isn't what happened. By their own admission after the experiments, the players indicated that they had no insight into what was happening. They "learned" how to react effectively to the specific behavior of the experimenter (all the while believing the equilibrium strategy the experimenter chose was foolish), but the subjects' ideas of what the game was about were as muddy at the end as they were at the start.

Apropos of "learning" how to play games, there is a theorem that is interesting enough to warrant a digression. Suppose that two people are playing a game repeatedly but do not have the skill to compute the minimax strategies. They both play randomly the first time, and subsequently they learn from their experience, in the following sense. They each assume that their opponent is playing a mixed strategy with probabilities proportional to the actual frequencies with which strategies have been chosen in the past. On the basis of this assumption, they play the pure strategy that will maximize their average return. Julia Robinson (1951) proved that, under these conditions, both strategies will approach the minimax strategies.

Of what practical significance are these experiments? Does it make any difference to a player to know that players often act irrationally? It didn't matter in games of perfect information such as chess. Does it matter now?

It is always possible to get the value of the game by playing minimax. That is so whether one's opponent is rational or not. The reason players are satisfied to get the value of the game is that they know a clever opponent can stop them from getting any more. But if players have reason to believe that their opponent will play badly, why not try to do better?

By using the minimax strategy, players will avoid doing

anything foolish, such as playing a dominated strategy. In our simplified poker game, players would always bet if they were dealt a 2; in the marketing game, they would never select three morning hours to advertise; and, in the military game, they would never send five divisions to one location. It should be emphasized, however, that the minimax strategy is essentially defensive, and when you use it, you often eliminate any chance of doing better. In Poe's odd-even game and in the game discussed by Rapoport, if a player adopts the minimax strategy, he or she obtains the value of the game; no less, but *no more.*

Suppose that, in the odd-even game, your opponent tends to choose even more often than he or she should. How can you exploit this if the game is played only once and you don't know how your opponent plays? Even if you know the game is so complex that he or she is not likely to analyze it correctly, how can you anticipate in what direction your opponent will err? In general, you may suspect that your opponent will pick the strategy that seems to yield the largest average payoff; but if you act on your suspicion, you become vulnerable to an opponent who is one step ahead of you and who is out to get rich by exploiting the exploiter. And if the game is played repeatedly and you manage to do better than you should, your opponent will eventually learn to protect himself or herself, so your advantage is unstable.

The weakest part of the theory is undoubtedly the assumption that a player should always act so as to maximize the average payoff. The justification for the assumption is that, in the long run, not only the average return but the actual return will be maximized. But if a game is played only once, are long-run considerations relevant? John Maynard Keynes (1972) made the point in *Monetary Reform:* "Now 'in the long run' this is probably true. . . . But this *long run* is a misleading guide to current affairs. *In the long run* we are all

dead." Often a strategy that maximizes the average return is not desirable, much less compelling. Is it irrational to opt for a sure million dollars rather than take an even chance of getting $10 million? This is not just an incidental objection; it goes to the heart of the matter. Let's take a closer look at this problem. In the game in figure 3.14, the payoffs are in dollars.

Figure 3.14

		PLAYER II	
		A	*B*
PLAYER I	*a*	$1 million	$1 million
	b	$10 million	$0
	c	$0	$10 million

The minimax strategy for player *I* is to play *b* and *c*, with a probability of one-half. Using this strategy, player *I* will get $10 million half the time and will get nothing half the time. But player *I* may prefer to have a sure million; in fact, rather than gamble, both players may prefer that player *I* win a sure million. This is one reason why cases are often settled out of court.

Things are not always what they seem. We started with what appeared to be a zero-sum game (it was a zero-sum game in dollars), but in fact the players had certain common interests. Yet the critical assumption in zero-sum games is that the *players have diametrically opposed interests.* And this must be so not only for the entries in the payoff matrix, which indicate the outcome when each player uses a pure strategy, but also for the various probability distributions that may be established when the players use mixed strategies.

This objection is almost fatal. It can be overcome only by introducing an entirely new concept, the concept of utility. (The word is old, but the idea is new.) The utility concept puts the theory on a firm foundation once again and in fact is

one of von Neumann and Morgenstern's most significant contributions to game theory. It will be discussed next.

Solutions to Problems

1. If your coin turns up 1, you should always bid. If your coin turns up 0, you should bid (bluff) one-eleventh of the time and pass the remaining ten-elevenths of the time.

   If your opponent's coin turns up 1, he or she should always see. If your opponent's coin turns up 0, he or she should see one-eleventh of the time and fold the remaining ten-elevenths of the time.

   You should win \$25/11 on average each time you play.

2. The police officers should never split up; they should go to the airport half the time and to the seaport the other half. The smugglers should split up four-fourteenths of the time; both should go to the airport five-fourteenths of the time and to the seaport five-fourteenths of the time. Fifty pounds of contraband will get through on average.

3. Ninety percent of the time the watchman should check the safe with \$90,000, (which is what one might suspect), but 90 percent of the time the burglar should try to rob the safe with \$10,000 (which is not so clear). On average the burgler should get \$9,000.

4. You should bid the slam with probability of .1 and your opponent should bid the slam with a probability of .9. In fact, the *less* likely the slam is to succeed the more likely your opponents will bid it. In this game your chance of winning the tournament is .91. (which is slightly better than the .9 you would get if you always bid game and your opponent always bid slam).

4

Utility Theory

Introductory Problems

Game theory is a tool for making decisions; but before you decide how to get what you want, you must first decide what you want. Sorting out what you want—the function of utility theory—is not always as easy as it sounds. Before reading the chapter, think about what you would do in each of the following situations.

1. When you arrive at a play, you find you have lost your pair of \$40 tickets. Would you spend another \$40 to buy a new pair of tickets or look elsewhere for your evening's entertainment?
2. Which would you prefer: a sure million dollars or an even chance at \$3 million?
3. A rich aunt dies and leaves you \$200; soon after a sporting uncle offers you the choice of another \$50 or a 25 percent chance of winning \$200. Which do you accept?
4. Suppose that 1 percent of the people of your age and health die in a given year. How much would you pay for \$100,000 worth of life insurance?
5. Would you prefer a sure \$10 or an even chance of \$30?

6. You arrive at a theater expecting to pay $40 for tickets and find you have lost $40 of your money. You still have more than $40 left—would you buy the tickets anyway?
7. A rich aunt (not the one in question 4) dies and leaves you $400. You are caught speeding and a whimsical judge offers you the following choice: You can pay a $150 fine, or you can pay $200 with a 75 percent probability and pay nothing with a 25 percent probability. What do you do?

The last chapter was primarily concerned with the minimax theorem—why it is needed, what it means, and why it is important. For the sake of clarity, I discussed the central ideas and avoided mentioning the difficulties. At the end of the chapter, however, it was suggested that all was not as it should be. Part of the foundation, and an absolutely essential part, is missing, and the theory is in fact firmly rooted in midair. One of von Neumann and Morgenstern's important contributions is the concept of utility, which makes the old solutions and strategies plausible once again. The utility function was fashioned precisely for this purpose; and in this chapter I will discuss utility functions and their role in game theory.

The heart of the difficulty is the ambiguity of the term "zero-sum." Taking the simplest view, a game is zero-sum if it satisfies a certain conservation law: a game is zero-sum if, during the course of the game, wealth is neither created nor destroyed. In this sense, ordinary parlor games are zero-sum.

But this definition won't do, for it is not the payoffs in money that are important. Throughout the discussion of the two-person, zero-sum game, it was assumed that each player was doing his or her best to hurt an opponent; if this assumption fails, the rest of the theory fails with it. What is needed to make the earlier argument plausible is the assurance that the players will indeed compete. A parent playing a card game with a child for pennies is playing a zero-sum game in

my original sense, yet the arguments of the last chapter do not apply. Of course, if we *assume,* as we did earlier, that a player's goal is to maximize the expected payoff in money, there is no problem, but this begs the question; it is the validity of this assumption that is in doubt.

As a matter of fact, there are many situations in which people *don't* act so as to maximize their expected winnings. This has nothing to do with the formal theory; it is an observation about life. The game used as an illustration at the end of the last chapter is only one example in which the assumption is incorrect; there are many others.

Generally, a person's wealth affects his or her attitude toward risk. A multimillion-dollar corporation will take a $50,000 gamble when a person with exactly that much capital will not. And this is so even when both perceive the situation in exactly the same way. Eli Schwartz and James Greenleaf (1978) suggest that this is one reason that the rich grow richer while the poor grow poorer. They constructed a model of a society in which everyone is equally wealthy initially and in which there was a sequence of bets that each could make. Some bets were riskier than others but their average return was higher as well. As chance made some members wealthier and others poorer, this difference was accentuated in a systematic way since the more affluent could afford the more speculative, higher-return bets. Schwartz and Greenleaf found in some computer runs that after five, twenty, and fifty trials the top tenth controlled 18 percent, 25 percent, and 50 percent of the total wealth, respectively.

Many games having a negative average payoff nevertheless attract a large number of players. For instance, people buy lottery tickets, play the "numbers game," bet at the race track, and gamble in Atlantic City, Las Vegas, and Reno. This isn't to say that people bet aimlessly without pattern or purpose; certain bets are generally preferred to others. In

studies of parimutuel betting at the race track, for instance, Griffith (1949) observed that bettors consistently gave more favorable odds on long-shots than experience warranted, while the favorites, at more attractive odds, were neglected. Just what it is that makes one bet more attractive than another is not always clear; what is clear is that the players are trying to do something other than maximize their average winnings.

It isn't just the gamblers who are willing to play games with unfavorable odds; such games are also played by conservatives who wish to avoid large swings. Commodity-futures markets are created because at planting farmers want to be sure prices don't drop too low by harvest. Also, almost every adult carries personal insurance of one sort or another, yet the "game" of insurance has a negative value, for the premiums cover not only the benefits of all policyholders but insurance company's overhead and commissions as well. Incidentally, insurance is just a lottery turned around. In both cases, the player puts in a small amount; lottery players have a small chance to win a fortune, and insurance policyholders avoid the small chance of having a catastrophe befall them.

The prevalence of insurance policies reflects willingness to pay a price for security; this aversion to risk has also been observed in the laboratory.

Harry Markowitz (1955) asked a group of middle-class people whether they would prefer to have a smaller amount of money with certainty or an even chance of getting ten times that much. The answers he received depended on the amount of money involved. When only a dollar was offered, all of them gambled for ten, but most settled for a thousand dollars rather than try for ten thousand, and all opted for a sure million dollars.

It should be emphasized that the failure to maximize average winnings is not simply an aberration caused by a lack of

insight. In 1959 Alvin Scodel, J. Sayer Minas, and Philburn Ratoosh, asked the subjects of an experiment to select one of several bets. There was no correlation whatever between the bets selected and the subject's intelligence. Graduate students in mathematics were included among the subjects, and they certainly could do the necessary calculations if they wished.

So if the theory is to be realistic—and if it is not realistic, it is nothing—we cannot assume that people are concerned only with their average winnings. In fact, we cannot make any general assumption about people's wants, because different people want different things. What is needed is a mechanism that relates the goals of a player, whatever they are, to the behavior that will enable him or her to reach these goals. In short, a theory of utility.

Before you can make sensible decisions in a game, both the goals of the players and the formal structure of the game must be taken into account. In this process of decision making, there is a natural division of labor. The game theorist, to paraphrase Lewis Carroll, must pick the proper road after being told the player's destination; and players don't have to know anything about game theory, but they have to know what they like—a slightly altered version of the popular bromide, "I don't know anything about art but I know what I like."

The problem is to find a way for players to convey their attitudes in a form that is useful to the decision maker. Statements such as "I detest getting caught in the rain" or "I love picnics" do not help someone decide whether to call off a picnic when the weather prediction is an even chance of rain. There is no hope of completely describing subjective feelings quantitatively, of course, but, using utility theory, it is possible to convey enough of these feelings (under certain conditions) to satisfy my present purpose.

Utility Functions: What They Are, How They Work

A utility function is simply a "quantification" of a person's preferences with respect to certain objects. Suppose I am concerned with three pieces of fruit: an orange, an apple, and a pear. The utility function first associates with each piece of fruit a number that reflects its attractiveness. If the pear was desired most and the apple least, the utility of the pear would be greatest and the apple's utility would be least.

The utility function not only assigns numbers to fruit; it assigns numbers to lotteries that have fruit as their prizes. A lottery in which there is a 50 percent chance of winning an apple and a 50 percent chance of winning a pear might be assigned a utility of 6. If the utilities of an apple, an orange, and a pear were 4, 6, and 8 respectively, the utilities would reflect the fact that the person was indifferent (had no preference) between a lottery ticket and an orange, that he or she preferred a pear to any other piece of fruit or to a lottery ticket, and that he or she preferred a lottery ticket or any other piece of fruit to an apple.

Also, utility functions assign numbers to all lotteries that have as prizes tickets to other lotteries; and each of the new lotteries may have as its prizes tickets to still other lotteries, so long as the ultimate prizes are pieces of fruit.

This is still not enough, however. Von Neumann and Morgenstern demand one more thing of *their* utility functions that make them ideally suited for their theory. The utility functions must be so arranged that the utility of any lottery is always equal to the weighted average utility of its prizes. If in a lottery there is a 50 percent chance of winning an apple (which has a utility of 4) and a 25 percent chance of winning either an orange or a pear (with utilities of 6 and 8, respectively), the utility of the lottery would necessarily be $5\frac{1}{2} = (.5)(4) + (.25)(6) + (.25)(8)$.

The Existence and Uniqueness of Utility Functions

It is easy enough to list conditions that you would like the utility functions to satisfy; it is quite another thing to find a utility function that does. Given a person of arbitrary tastes, is it always possible to find a utility function that reflects them? Can these tastes be reflected by two different utility functions?

Taking the second question first, the answer is yes. Once a utility function is established, another can be obtained from it by simply doubling the utility of everything; and still another, by adding one to the utility of everything. In fact, if we take any utility function and multiply the utility of everything by any positive number (or add the same number to every utility), a new utility function is obtained that works just as well.

It is also possible that there is no utility function that will do; this would be the case if the person's tastes lacked "internal consistency." Suppose a person, when given a choice between an apple and an orange, prefers the apple. That would mean that the utility of an apple would have to be greater than the utility of an orange. If, between an orange and a pear, the person prefers the orange, the utility of an orange would have to be greater than that of a pear. Since the utilities are ordinary numbers, it follows that the utility of an apple is necessarily greater than the utility of a pear. If, when given the choice between an apple and a pear, the person chooses the pear, establishing a utility function would be hopeless. There is no way one can assign numbers to the pieces of fruit that will simultaneously reflect all three of these preferences. There is a technical name for what I have just described: the player's preferences are said to be *intransitive.*

If the preferences of a player are sufficiently consistent—that is, if they satisfy certain requirements—the preferences can be expressed concisely in the form of a utility function.

GAME THEORY

Six Conditions That Guarantee the Existence of a Utility Function

If a player's preferences are to be expressed by a utility function, these preferences must be consistent; that is, they must satisfy certain conditions. These conditions may be expressed in several more or less equivalent ways; I have used the formulation suggested by Duncan R. Luce and Howard Raiffa (1957). For convenience, the word "object" will represent either a piece of fruit or a lottery.

1. *Everything is comparable.* Given any two objects, the player must prefer one to the other or be indifferent to both; no two objects are incomparable.
2. *Preference and indifference are transitive.* Suppose A, B, and C are three different objects. If the player prefers A to B and B to C, A will be preferred to C. If the player is indifferent between A and B, and B and C, he or she will be indifferent between A and C.
3. *A player is indifferent when equivalent prizes are substituted in a lottery.* Suppose in a lottery one prize is substituted for another but the lottery is left otherwise unchanged. If the player is indifferent between the old and new prizes, he or she will be indifferent between the lotteries. If the player prefers one prize to the other, he or she will prefer the lottery that offers the preferred prize.
4. *A player will always gamble if the odds are good enough.* Suppose that, of three objects, A is preferred to B and B is preferred to C. Consider the lottery in which there is a probability p of getting A and a probability of $1 - p$ of getting C. Notice that if p is zero the lottery is equivalent to C, and if p is 1, the lottery is equivalent to A. In the first case the lottery is preferable to B, while in the second case B is preferable to the lottery. According to this condition, that there is a value p between zero and 1 that will make the player indifferent between B and the lottery.
5. *The more likely the preferred prize, the better the lottery.* In lotteries *I* and *II* there are two possible prizes: objects A and B. In lottery *I*, the chance of getting A is p; in lottery *II*, the chance of getting A is q. A is preferred to B. This condition requires that if p is bigger than q, lottery *I* be preferred to lottery *II*; and, conversely, if lottery *I* is preferred to lottery *II*, p is greater than q.
6. *Players are indifferent to gambling.* A player's attitude toward a compound lottery—a lottery in which the prizes may be tickets to other lotteries—is dependent only on the ultimate prizes and the chance of getting them as determined by the laws of probability; the actual gambling mechanism is irrelevant.

Henceforth it will be supposed that the preferences of each player are expressed by a utility function and that the payoffs are given in "utiles"—the units in which utility functions are expressed. When I speak of a zero-sum game now, I mean a game in which the sum of all the payoffs (in utiles) is always zero. When two players are involved in a zero-sum game in this new sense, their interests must necessarily be opposed.

The advantage of this new definition of zero-sum is that, if applied to the new zero-sum games, it justifies the work earlier on two-person, zero-sum games. But, admittedly, there is a corresponding disadvantage in the new approach. This might have been anticipated as a consequence of our quasi-conservation law: games that are easy to analyze don't come up very often. The difficulty is that this new type of zero-sum game is rare and hard to recognize. Before I changed the definition, you only had to know the rules of a game such as poker to see immediately that it was zero-sum. Now the problem of recognition is much more subtle and involves subjective factors such as the players' attitude toward risk. This makes the theory much more troublesome to apply.

GAME THEORY

Some Potential Pitfalls

Utility theory is easily misunderstood. The reason for this is partly historical: "utilities" have been in existence for some time, but the term has not always been used in a consistent or clear way. In an article called "Deterrence and Power," Glenn Snyder (1960) gave a number of good examples of how one can go wrong, as follows:

> The figures which follow [i.e., the payoff matrix] are based on the assumptions that both sides are able to translate all relevant values into a "common denominator" "utility," that they can and do estimate probabilities for each other's moves, and that they act according to the principle of "mathematical expectations." The latter states that the "expected value" of any decision or act is the sum of the expected values for all possible outcomes, the expected value for each outcome being determined by multiplying its value to the decision-making unit times the probability that it will occur. To act "rationally" according to this criterion means simply to choose from among the available courses of action the one which promises to maximize expected value (or minimize expected cost) over the long run. There are reasons—apart from the practical difficulty of assigning numerical values to the elements involved—why the mathematical expectation criterion is not entirely suitable as a guide to rationality in deterrence and national security policy. However, it is useful as a first approximation, the necessary qualifications—having to do with the problem of uncertainty and the disutility of large losses—will be made presently. (p. 168)

Snyder starts out all right; you do have to "translate all relevant values into a 'common denominator' 'utility.'" That this is possible must be assumed, but that is *all* you have to assume. Snyder is also willing to assume that "they act according to the principle of 'mathematical expectations.'" But why should a person, or a country, want to maximize its

utility, any more than it would want to maximize its average winnings in money?

Both the last question and Snyder's assumption put the cart before the horse. The person's wants come first; the utility function, if it exists, comes second. The person is *not* trying to maximize his or her utility—average or otherwise; very likely, the person doesn't even know such a thing exists. Players may act *as though* they were maximizing their utility function, but this is not because they mean to but because of the judicious way the utility function was established. What happens, at least in theory, is that the preferences of the players are observed and then a utility function is established that the players *seem* to be maximizing.

Snyder's statement "To act 'rationally' according to this criterion means simply to choose from among the available courses of action the one which promises to maximize expected value (or minimize expected cost) over the long run" can be criticized on two counts. For one thing, as with the earlier statement, the problem is put the wrong way around. You don't maximize your utility because you're rational; your preferences are observed *first* and *then* the utility function is established. Moreover, "over the long run" has nothing to do with it. Utilities are assigned to one-shot gambles. If a person takes the view that life is a series of gambles in which good and bad luck cancel each other, she may decide to maximize her average return in dollars; and this is all very well. But if she prefers less risky alternatives, that's all right too, even if it lowers her expected return. Utility theory will accommodate either attitude. In the deterrence "game," gambles are taken without any assurance that they will recur; in fact, there is good reason to believe they will not.

And, finally, Snyder's reservations about the "disutility of large losses" entirely misses the point. Either a utility func-

tion exists or it doesn't. If the utility function associates with large losses numbers that are inaccurate, then it isn't a utility function at all. A utility function, if it does anything, must reflect a person's preferences accurately; it serves no other purpose.

Constructing a Utility Function

Granted that there is a need for a utility function, there still remains the problem of actually determining it. Roughly speaking, this is how it can be done. People are asked to make many simple choices between two things; it may be pieces of fruit, or lottery tickets. People might be asked, for example, whether they would prefer a ticket to a lottery in which there is a three-fourths chance of winning a pear and a one-fourth chance of winning an apple, or the certainty of getting an orange. On each question they must indicate a preference for one of the alternatives or state that they are indifferent between both. It is possible, on the basis of these simple choices, to establish a single utility function that assigns a number to every piece of fruit and every possible lottery—reflecting all the preferences of an individual simultaneously, providing players' preferences are consistent. In effect, the fruit and lotteries are put on a single dimension in which every piece of fruit and every lottery ticket can be compared simultaneously.

It should be understood that in the process of establishing a utility function, nothing new was added; the grand final ordering is implicit in the simple choices made earlier. But the practical advantage of having a concise utility function rather than a great many individual preferences is enormous.

Are People's Preference Patterns Really Consistent?

So far I have assumed that people for whom utility functions are to be established have consistent preferences; that is, preferences that satisfy the six consistency conditions listed earlier. On the face of it, these conditions seem intuitively appealing, and one would think that most people would accept them with few if any objections. It has even been suggested that these conditions—or, at any rate, some of them—should be used as a *definition* of rationality in decision making. But it turns out that people often have, or seem to have, inconsistent attitudes that preclude the construction of a utility function. Let's look at some of the things that can go wrong.

A fair amount of effort has been expended to study betting behavior experimentally and, in particular, to see whether it is consistent. In a number of experiments, Ward Edwards (1953, 1954 *a*, *b*, *c*, *d*) compared the bets that subjects chose with certain other variables, such as the average amount the subject could win, the state of the subject's fortunes when the choice was made, the actual numerical probabilities involved, and so forth. He found, among other things, that certain subjects tended to make bets in which the probabilities involved were one-half, one-half rather than one-quarter, three-quarters, even though the average return was the same in both cases. It is easy to show that a person cannot consistently prefer the one type of bet to the other and still satisfy the six conditions necessary for consistency.

Another source of difficulty is the variation in people's preferences over a period of time. This might not appear to present too serious a problem, since variations would occur gradually; but that is not the case. The interaction between

what happens in the game and the attitudes of the players cannot be neglected at any time. A worker who has been out on strike for a long time looks at an offer with different eyes than a worker just starting to negotiate; and, during the process of negotiation, flexible attitudes may harden and what were once acceptable possibilities may become unacceptable.

Anyone who has watched people gambling will no doubt have observed the same thing. In a poker game, for example, as time passes, the bets grow larger: apparently the amount that a player considers an acceptable risk increases as the game proceeds. The same phenomenon has been observed experimentally by comparing the size of bets made in the early and the late races at the track. A subjective but graphic picture of what happens is given by Dostoevsky (1915) in his novel *The Gambler*:

> As for me, I lost every farthing quickly. I staked straight off twenty friedrichs d'or on even and won, staked again and again won, and went on like that two or three times. I imagine I must have had about four hundred friedrichs d'or in my hands in about five minutes. At that point I ought to have gone away, but a strange sensation rose up in me, a sort of defiance of faith, a desire to challenge, to put out my tongue at it. I laid down the largest stake allowed—four thousand gulden—and lost it. Then getting hot, I pulled out all I had left, staked it on the same number, and lost again, after which I walked away from the table as though I were stunned, I could not even tell what happened to me. . . . (p. 168)

There are many other problems as well. Experiments have shown that decisions often depend on seemingly irrelevant variables. People make one type of bet when playing for money and another when playing for wooden chips that are worth money. They bet one way when others are present and another way when they're alone. Their history—the success they've had in the game so far—influences their attitude

toward risk. People select inconsistently; they pick one of two bets one time, and the other the next. And their preferences are intransitive: they prefer A when offered a choice between A and B, and B when offered a choice between B and C. And when offered a choice between A and C, they prefer C.

Note how the context of a problem, as described by Leonard J. Savage (1954), can influence decision making: "A man buying a car for $2,134.56 is tempted to order it with a radio installed, which will bring the total price to $2,228.41, feeling that the difference is trifling. But when he reflects that, if he already had the car, he would certainly not spend $93.85 for a radio for it, he realizes that he had made an error" (p. 103).

Recently serious doubts were raised about the foundations of utility theory. In an experiment conducted by Daniel Kahneman and Amos Tversky (1982), it was found that most subjects made different decisions in two identical or equivalent situations when these situations were described in different ways. The general rule was this: people who feel they have won something generally try to conserve their winnings by avoiding risks. In an identical situation, the same people who perceive that they have just lost something will take risks they considered unacceptable before, to make themselves whole. (See "Solutions to Problems," p. 73, for more details on this experiment.)

In the same experiment, subjects were told a jacket cost $125 and a calculator $15. But they were offered a chance to get the calculator for $10 if they drove twenty minutes to another store. Most accepted. But when the prices were switched and $5 was saved on the $125 item instead of the $15 item, most subjects declined to travel; presumably a decrease of 4 percent was considered less significant than a decrease of 33⅓ percent. But in fact this is another variation of Savage's paradox, because in both cases you save $5 by traveling 20 minutes.

Despite all the difficulties—despite the seeming irrationality of people's behavior, despite the inconsistencies—utility functions have been established successfully. In one case experimenters observed the bets that subjects chose when they were offered simple alternatives. On the basis of these observations, they were able to predict what the subjects would do when presented with much more complicated decisions. How was this possible, in view of all the objections mentioned?

Not all the objections can be explained away, of course, but the problem is not as serious as it may seem. Take the case of intransitive preferences, for instance. Some experimenters feel that true intransitivities rarely arise. What happens, they say, is that people are forced to choose between alternatives to which they are indifferent. The answer, then, is dependent on momentary whim. If people have difficulty deciding whether to buy a car for $7,000, it should not be surprising if one moment they refuse to buy at a dollar less and the next moment they agree to buy at a dollar more. In several experiments it was observed that subjects who make intransitive choices don't make them consistently. In one experiment, for instance, subjects were asked to list certain crimes in the order of their seriousness and to determine the severity with which guilty parties should be punished. While occasional intransitivities turned up, when the task was repeated they almost invariably disappeared.

It is likely that much "irrational behavior" can be avoided by identifying and controlling the significant variables. If, for example, people bet more aggressively when they are with others than when they are alone (and they do), this must be taken into account. Also, the subjects should have some experience in decision making, so they can take into consideration the consequences of their actions. In Savage's example, this is precisely what was needed.

Finally, people must be properly motivated. In some experiments subjects stated that they had varied from what they believed to be the most prudent course, just to keep things exciting. This is a serious problem, since as a rule experimenters have limited resources and cannot pay enough to motivate the players to the degree they would like. People betting in a "hot, humid, exciting, and highly disordered carnival setting," as one experimenter put it, behaved quite differently from bettors in an artificial, insulated, experimental environment.

Solutions to Problems

1, 3, 6, and 7. The point of these questions is that people do not always make decisions based on their situations but rather on how their situations arose or were described. You are more likely to buy tickets in the situation of question 1 than in question 6 if you are like the subjects of an experiment by Kahneman and Tversky (1982). Also, most subjects declined the gamble offered in question 3 but accepted the gamble in question 7 although the two cases were identical: a sure $250 versus a 25 percent chance at $400 and a 75 percent chance of $200. This apparent inconsistency raises doubts about the foundations of game theory.

When people are faced with a loss, they seem to seek out risks—they want to gamble to remain whole. The loser in a poker game who wants to play double or nothing and the holder of a losing stock waiting to break even are two examples of this kind of behavior. People who are winners often want to avoid risk—witness the winning poker player who wants to go home. This seems to account for the difference in behavior because the objective situations are identical.

2, 4, and 5. Game theory does not tell you whether you should gamble or not—it tells you how to achieve what you

want. There is nothing irrational about preferring a sure million dollars in question 2 and gambling for $30 in question 5—in fact, most people would do the same. The "fair" price of insurance is $1,000 in question 4 (this would mean you would break even if you took this bet many times), but there is no reason you can't pay more if you want insurance and no reason you can't decline a lower premium if you don't want insurance.

5

The Two-Person, Non-Zero-Sum Game

Introductory Problems

The games I have been discussing until now have been zero-sum—your gain was your opponent's loss. The non-zero-sum games that we encounter are at once more commonplace, more interesting, and more difficult to analyze. There are theories that favor one strategy over another but the arguments are less persuasive; the ultimate outcome can be conjectured about but the predictions are less compelling. The problems in this and the final two chapters have more to do with ideas than numbers. For this reason prethinking about the introductory ideas—before the text shapes your ideas—will be particularly rewarding.

1. You may assume that the matrix entries in figure 5.1 are in dollars. With each pair of strategies there are *two* payoffs: one for you (shown first), and one for your partner. If you choose *B* and your partner, *A*, you would get $5 and your partner would get nothing. Communication between players is not allowed.

Figure 5.1

		YOUR PARTNER: *A*	YOUR PARTNER: *B*
YOU	*A*	(3, 3)	(0, 5)
	B	(5, 0)	(1, 1)

   a. What strategy would you choose if this game were played only once?
   b. Would you play any differently if the game was to be repeated 200 times?
   c. Some time ago a number of people were asked to write a computer program describing exactly how to play on each of 200 repetitions of the game; your play on the fiftieth round would presumably depend on the plays on earlier rounds. In a tournament, each program played round-robin against every other program. (In addition, one program was played randomly.) Before deciding what overall strategy you would choose, think about the implications of each strategy.*
2. Suppose that you and another player are engaged in a game in which preplay communication is allowed.
   a. Suppose, further, that the game is changed so that you are prevented from using some of your strategies; can this possibly be to your advantage?
   b. Communications are cut off; can this be to your advantage? (Assume that you could make binding agreements when communication was available.)
   c. Suppose the rules in the original game required you and the other player to make simultaneous choices of strategy. However, the rules are changed so that your decision must be made first and the other player chooses after seeing what you have done; can this be to your advantage?

* This is a particularly important game and really worth some thought—take the time to sort it out.

   d. Will it be to your advantage or disadvantage if the other player learns what your utility function is?
3. A customer offers to buy a lamp from you, an antique dealer, for $5,000 within twenty-four hours. A wholesaler, who knows of the customer's offer, obtains the lamp for $3,000. At the last minute you get a note from the wholesaler offering the lamp for sale. You have only enough time to accept the offer and take your profit or turn it down and lose the sale. Would you accept the offer if:
   a. The selling price was (i) $4,000? (ii) $4,500? (iii) $4,950?
   b. Would it make any difference if you found out that the wholesaler mistakenly believed that the customer offered $7,000 rather than $5,000?
4. A speed limit was recently imposed in a community and the town council must now decide how strictly it should be enforced. By quantifying various costs and benefits to the community and the drivers—the time saved by speeding, the danger to the driver, the danger to the public in general, the penalties to the speeding driver, the cost of enforcement—the town game theorist arrives at the payoff matrix depicted in figure 5.2.

Figure 5.2

		COMMUNITY: ENFORCE THE LAW	COMMUNITY: DO NOT ENFORCE THE LAW
DRIVER	SPEED	(−190, −25)	(10, −5)
	DO NOT SPEED	(0, −20)	(0, 0)

   The matrix reflects the payoffs if the law is enforced 100 percent of the time. Is it "cheaper" for the community to enforce or ignore the law? Is there a third alternative?
5. You and a stranger each pick a number from 1 to 10. A philanthropic organization pays the one who picked the *lower* number as many thousands of dollars as the number picked. (If you both pick the same number, the winner is determined by a coin toss.)

a. What number should you choose if the game is played: (i) once? (ii) fifty times?
b. The game is to be played fifty times. On the first two plays, the stranger chooses 10 and 9. What do you make of this?

6. This last example is a kind of psychological test that measures your, and a hypothetical partner's, joint aspirations. Imagine that you and your partner are each given an algebraic expression after you have been placed in separate rooms to avoid communication. In example (a) that expression is $2X + 5Y - 100$. You choose a number X and your partner chooses a number Y. If the numbers you choose jointly make the expression positive, you get nothing; if they make it negative or zero, you get $\$X$ and your partner gets $\$Y$. So if you choose $X = 10$ and your partner chooses $Y = 2$, you get \$10 and your partner \$2 since $(2)(10) + (5)(2) - 100 = -70$, which is negative. You both get nothing if you choose 40 and your partner chooses 5, since $(2)(40) + (5)(5) - 100 = 5$, which is positive. You want to make X as large as possible so long as it isn't so large as to reduce your payoff to zero.

   To help visualize the problem, I have drawn the graph of the equation $2X + 5Y - 100 = 0$ (and all the other graphs as well). (See figure 5.3 *A–H.*) Interpreted geometrically, if the point (X, Y) is below the graph or on it, you and your partner get $\$X$ and $\$Y$, respectively. If it is above the graph, you both get nothing.

   If you are going to act intelligently, you should know something about your partner; so assume, as usual, that he or she is as bright as you are and has the same general attitude toward money that you do.

THE theory of two-person, zero-sum games is unusual in that it enables you to find solutions—solutions, moreover, that are universally accepted. In this respect, zero-sum games are unlike the actual problems that arise in everyday life, which generally do not lend themselves to straightforward answers. For non-zero-sum games we cannot do nearly as

Figure 5.3

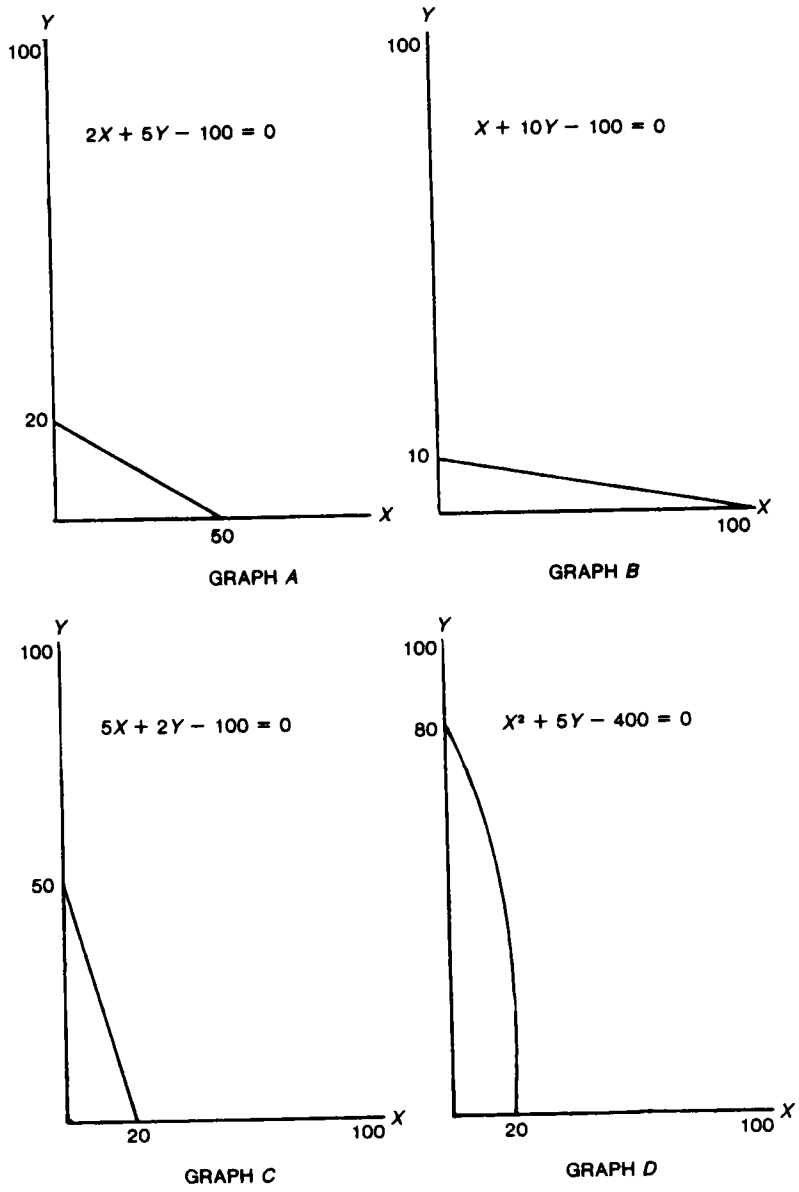

Figure 5.3 (*cont.*)

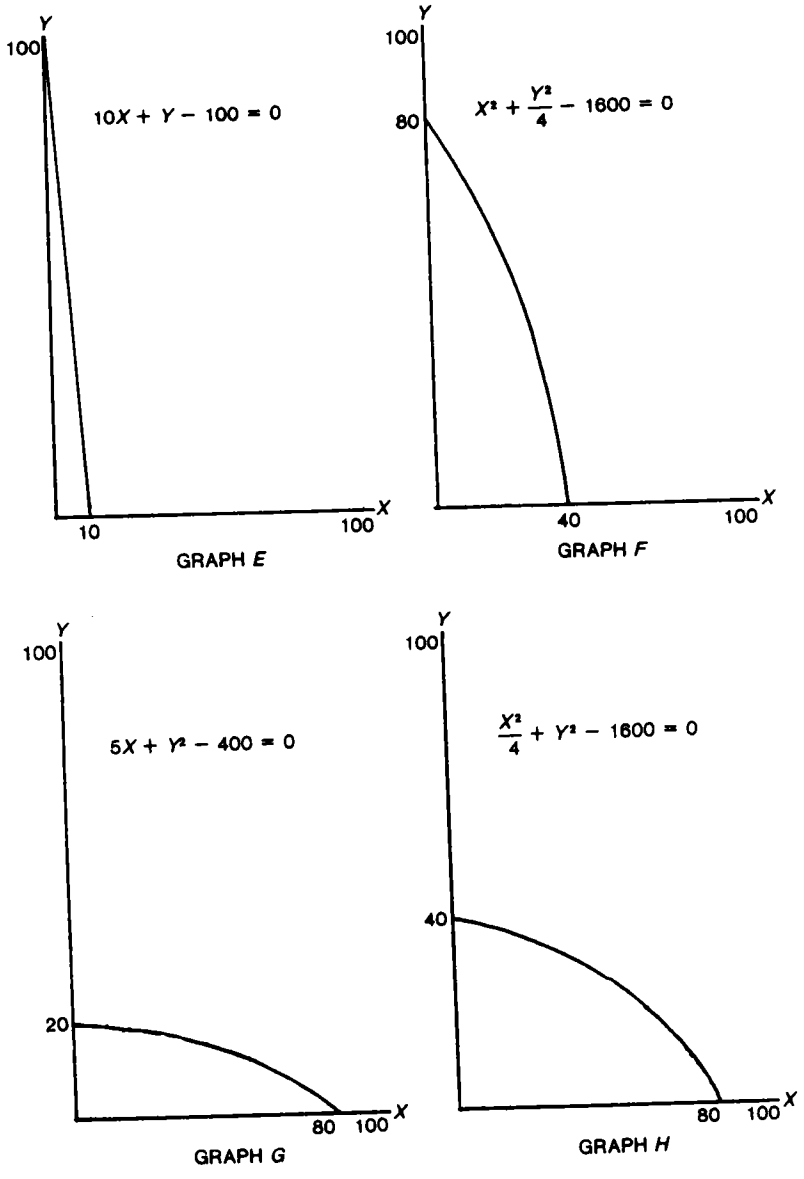

well. In most games of any complexity, there is no universally accepted solution; that is, there is no single strategy that is clearly preferable to the others, nor is there a single, clear-cut, predictable outcome. As a rule, we will have to be content with something less than the unequivocal solutions we obtained for zero-sum games.

For convenience, let's regard all two-person games as lying in a continuum, with the zero-sum games at one extreme. In a two-person game, there are generally both competitive and cooperative elements: the interests of the players are opposed in some respects and complementary in others. In the zero-sum game, the players have no common interests. In the completely cooperative game at the other extreme, the players have nothing but common interests. A pilot of an airplane and the operator in the control tower are engaged in a cooperative game in which they share a single common goal, a safe landing. Two sailboats maneuvering to avoid a collision, and two partners on a dance floor, are also playing a cooperative game. The problem in such a game is easy to solve, at least conceptually: it consists of coordinating the efforts of the two players efficiently (by means of dancing lessons, for instance).

The rest of the two-person games, and the ones with which this chapter will be primarily concerned, fall between these two extremes. Games with both cooperative and competitive elements are generally more complex, more interesting, and encountered more frequently in everyday life than pure competitive and cooperative games. Some examples of these games are: an automobile salesperson negotiating with a customer (both want to consummate the sale, but they differ on the price), two banks discussing a merger, two competing stores, and so forth. In each of these games the players have mixed motives. Also, there are many situations where the parties *seem* to have no common interests but really do.

Two nations at war may still honor a cease-fire, not use poison gas, and refrain from using nuclear weapons. In fact, zero-*sum* games are almost always an approximation to reality—an ideal that is never quite realized in practice. In chapter 4, for example, if the examples are varied even slightly, they cease to be zero-sum. In the marketing problem, it was assumed that total consumption was constant. But if, when the advertising schedule was changed, total consumption also changed (as in practice it would), the game would be non-zero-sum. And even if the assumption were valid but the firms conspired together to reduce the amount of advertising (and thus lower costs), the game would also be non-zero-sum.

Analyzing a Two-Person, Non-Zero-Sum Game

The easiest way to get some insight into the non-zero-sum game is to try to analyze one in the same way that the zero-sum game was analyzed. Suppose we start with the matrix shown in figure 5.4. Notice that in non-zero-sum games it is necessary to give the payoffs for *both* players, since the payoff of one player can no longer be deduced from the payoff of the other, as it can in zero-sum games. The first number in the parentheses is the payoff to player *I*, and the second number is the payoff to player *II*.

Figure 5.4

		PLAYER II	
		A	*B*
PLAYER I	*a*	(0, 0)	(10, 5)
	b	(5, 10)	(0, 0)

It is obvious what the players' common interests is to avoid the zero payoffs. But there still remains the problem of determining who will get 5 and who will get 10. One way to attack the problem, a way that was successful before, is to look at the game from the point of view of one of the players—say, player *I*. As player *I* sees it, the game has the matrix shown in figure 5.5. Since the payoffs to player *II* don't concern that player directly, they are omitted.

Figure 5.5

		PLAYER II	
		A	*B*
PLAYER I	*a*	0	10
	b	5	0

Ultimately, player *I* must decide what she wants from this game and how she should go about getting it. A good way to start is by determining what she can get without any help from player *II;* this, at any rate, is the *least* player *I* should be willing to settle for.

Using the techniques of the zero-sum game, player *I* will find that, by playing strategy *A* one-third of the time she can get an average return of 10/3. If player *II* plays strategy *A* two-thirds of the time, that's all player *I* can get.

If player *II* makes the same sort of calculation, using his own payoff matrix and ignoring the payoff to player *I*, he too can get 10/3 by playing strategy *A* one-third of the time. If player *I* plays strategy *a* two-thirds of the time, player *II* will not get any more.

So far, so good. It looks as if we are in much the same position as we were in the zero-sum game: each player can get 10/3; each player can be stopped from getting any more. Why not, then, call this the solution?

The trouble with this "solution" is that the payoffs are too low. If the players can manage to get together on either one of the non-zero payoffs, they will *both* do better. The argu-

ment used to support the (10/3, 10/3) payoff, which was sound enough in the zero-sum game, fails here. *Although each player can prevent the other from getting more than 10/3, there is no reason why he or she should.* It is no longer true that a player can get rich only by keeping an opponent poor. Although it is true that the players cannot be prevented from getting 10/3, they are foolish if they don't aspire to something better.

But how to go about doing better? Suppose player *I* anticipates everything that I have said so far and guesses that players *II* will play strategy *A* one-third of the time, to ensure a payoff of 10/3. If player *I* switches to the pure strategy *a*, player *II* will still get his 10/3, but player *I* will now get 20/3: twice as much as before. The trouble with this argument is that player *II* may get ambitious also. This player may also decide to double his return and, anticipating conservative play by player *I*, play the pure strategy *A*. If they both do this simultaneously, they will both get nothing. As in the zero-sum game, this sort of circular reasoning does not take you very far.

Up to now, my main purpose has been to give some idea of the kind of problems that may arise. To this end, and also for future reference, let us look at some examples.

The Gasoline Price War

Two competing gas stations buy gas at 20 cents a gallon and, together, sell a thousand gallons a day. They have each been selling gas at 25 cents a gallon and dividing the market evenly. One of the owners is thinking of undercutting the other. The prices, which are quoted in whole numbers, are

set independently by each station in the morning and remain fixed for the day. There is no question of customer loyalty, so a lower price at either station attracts virtually all the sales. What price should each station owner set? If one owner knows the other is going out of business the next day, should that make a difference? What if the owner is going out of business ten days later? What if the owner is going out of business at some indefinite time in the future?

For each set of prices set by the gas stations, corresponding profits are entered in the matrix in figure 5.6. If gas station *I* charges 24 cents and gas station *II* charges 23 cents, gas station *II* captures the entire market; its owner makes 3 cents profit per gallon on a thousand gallons, or $30 profit. Gas station *I* makes no profit at all.

Figure 5.6

GAS STATION *I*'S PRICE PER GALLON	GAS STATION *II*'S PRICE PER GALLON: 25¢	24¢	23¢	22¢	21¢
25¢	(25, 25)	(0, 40)	(0, 30)	(0, 20)	(0, 10)
24¢	(40, 0)	(20, 20)	(0, 30)	(0, 20)	(0, 10)
23¢	(30, 0)	(30, 0)	(15, 15)	(0, 20)	(0, 10)
22¢	(20, 0)	(20, 0)	(20, 0)	(10, 10)	(0, 10)
21¢	(10, 0)	(10, 0)	(10, 0)	(10, 0)	(5, 5)

Let us start by trying to answer the easiest question: What price should a gas station set if its competitor is going out of business the next day? The first thing to note is that one should *never* charge less than 21 cents or more than 25 cents. When the price drops below 21 cents, there is no profit, and when it rises above 26 cents, you lose your market. Note also that you should never charge 25 cents, since this price is dominated by 24 cents. If your competitor charges either 24 or 25 cents, you do better by charging 24 cents, and if your

competitor charges something else, you don't do any worse. A gas station should not charge 25 cents, then, and the owner may reasonably assume that a competitor won't either.

Once 25 cents is ruled out as a selling price for *either* gas station, you can, by the same reasoning, rule out 24 cents as a selling price, because it is dominated by a price of 23 cents. As a matter of fact, you can successively rule out every price but 21 cents.

There is a paradox here. Starting from the status quo, where the stations charged 25 cents a gallon and received $25 profit apiece, they managed, by reasoning logically, to get to a position where they charge 21 cents and receive only $5 profit apiece.

If one of the gas stations is known to be going out of business in ten days rather than in one day, it's a bit more complicated. The old argument is all right as far as it goes, but now we are concerned with the profits not only of the next day but of the following nine as well. And if a player drops his or her price one day, that player can be sure that the competitor will do the same on the next. It turns out, however, that even here an apparently airtight argument can be made for charging 21 cents. It is only when the competition continues indefinitely that the argument finally loses force.

This game is a variation of the prisoner's dilemma; I will say more about it later.

A Political Example

The state legislature is about to vote on two bills that authorize the construction of new roads in cities *A* and *B*. If the two cities join forces, they can muster enough political

power to pass the bills, but neither city can do it alone. If a bill is passed, it will cost the taxpayers of both cities a million dollars and the city in which the roads are constructed will gain $10 million. The legislators vote on both bills simultaneously and secretly; each legislator must act on each bill without knowing what anyone else has done. How should the legislators from cities *A* and *B* vote?

The payoff matrix for this game is shown in figure 5.7. Since a city always supports its own bond issue, the legislators have only two strategies: supporting, or not supporting, the sister city.

Figure 5.7

		CITY *B*: SUPPORT *A*'S BOND ISSUE	CITY *B*: WITHHOLD SUPPORT FROM *A*'S BOND ISSUE
CITY *A*: SUPPORT *B*'S BOND ISSUE		(8, 8)	(−1, 9)
CITY *A*: WITHHOLD SUPPORT FROM *B*'S BOND ISSUE		(9, −1)	(0, 0)

The entries in the matrix are in millions of dollars. As an example of how they were computed, suppose city *A* supports city *B*'s bond issue, but *B* doesn't support *A*'s. Then one bond issue passes, each city pays a million dollars, *B* gets $10 million, and *A* gets nothing. The net effect is that *A* loses a million dollars and *B* gets $9 million.

As before, the smartest thing seems to vote against the other city's bond issue. And, just as before, if both cities follow this strategy, they will both get nothing, rather than getting the $8 million apiece they could have gotten otherwise.

The Battle of the Sexes

Another interesting two-person, non-zero-sum game—the battle of the sexes—is described by R. Duncan Luce and Howard Raiffa (1957). A husband and his wife have decided that they will go either to the ballet or to the prize fights that evening. Both prefer to go together rather than going alone. While the husband would most prefer to go with his wife to the fights, he would prefer to go with her to the ballet rather than go to the fights alone. Similarly, the wife's first choice is that they go together to the ballet, but she too would prefer that they go to the fights together rather than go her own way alone. The matrix representing this game is shown in figure 5.8. The payoffs reflect the order of the players' preferences. This is essentially the same game I talked about at the opening of this chapter.

Figure 5.8

		HUSBAND: FIGHTS	HUSBAND: BALLET
WIFE	FIGHTS	(2, 3)	(1, 1)
	BALLET	(1, 1)	(3, 2)

A Business Partnership

A builder and an architect are jointly offered a contract to design and construct a building. They are offered a single lump sum as payment and given a week to decide whether they will accept. The approval of both is required, so they must decide how the lump sum will be divided. The architect writes to the builder, suggesting that the profits be divided equally. The builder replies, committing himself to the contract in writing, provided he receives 60 percent of the

profits. He informs the architect, moreover, that he is going on a two-week vacation, that he cannot be reached, and it is up to the architect to accept the contract on these terms or reject it entirely.

The architect feels she is being exploited. She feels that her services are as valuable as the builder's and that each should receive half the profits. On the other hand, the probable profit is large, and under different circumstances, she would consider 40 percent of such a total an acceptable fee. Should she accept the builder's offer? Note that the builder has no more options; only the architect has a choice—to accept or reject the contract. This is reflected in the payoff matrix shown in figure 5.9.

Figure 5.9

		BUILDER
ARCHITECT	ACCEPT OFFER	(40%, 60%)
	REJECT OFFER	(0, 0)

A Marketing Example

Each of two competing firms is about to buy an hour of television time to advertise its products. They can advertise either in the morning or in the evening. The television audience is divided into two groups. Forty percent of the people watch during the morning hours, and the remaining 60 percent watch in the evening; there is no overlap between the groups. If both companies advertise during the same period, each will sell to 30 percent of the audience watching and make no sales to the audience not watching. If they advertise during different periods, each will sell to 50 percent of the audience watching. When should they advertise? Should

they consult before making a decision? The payoff matrix for the game is shown in figure 5.10.

Figure 5.10

		FIRM II: ADVERTISE IN MORNING	FIRM II: ADVERTISE IN EVENING
FIRM I	ADVERTISE IN MORNING	(12, 12)	(20, 30)
	ADVERTISE IN EVENING	(30, 20)	(18, 18)

The matrix entries represent the percentage of the total audience that each firm captures. If firm *I* picks an evening hour and firm *II* picks a morning hour, firm *I* sells to half of 60 percent, or 30 percent, and firm *II* sells to half of 40 percent, or 20 percent.

Some Complications

Two-person, zero-sum games come up in many different contexts, but they always have the same basic structure. By looking at the payoff matrix, you can pretty well tell "the whole story." This is not the case in non-zero-sum games. Besides the payoff matrix, there are many "rules of the game" that markedly affect the character of the game, and these rules must be spelled out before you can talk about the game intelligently. It is impossible to say very much on the basis of the payoff matrix alone.

To see this more clearly, go back to the first example discussed on page 82—in which the payoffs were (0, 0), (10, 5), and (5, 10). There were several details that I deliberately omitted but which are obviously important. Can the players consult beforehand and agree to their respective strategies in

advance? Are these agreements binding? That is, will the referee or whoever enforces the rules insist that the agreement be carried out, or does the agreement have moral force only? Is it possible for the players, after one has received 5 and the other 10, to make payments to each other so they both receive 7 1/2? (In some games, they can; in others, they can't.) In the gasoline price war, is it feasible (legal) for the two businesses to conspire to fix prices?

It is to be expected that these (and other) factors will have a strong influence on the outcome of the game; but their actual effect is often different from what one would expect.

After you study the two-person, zero-sum game, certain aspects of the non-zero-sum game seem like something out of *Alice in Wonderland.* Many "obvious" truths—fixed stars in the firmament of the zero-sum game—are no longer valid. Look at a few:

1. You would think that the ability to communicate could never work to a player's disadvantage. After all, even if players have the option to communicate, they can refuse to exercise it and so place themselves in the same situation as if facilities for communication were nonexistent; or so it would seem. The facts are otherwise, however. The *inability* to communicate may well work to one player's advantage, and this advantage is lost if there is a way to communicate, even though no actual communication occurs. (In the zero-sum game, this doesn't come up. There the ability to communicate is neither an advantage nor a disadvantage, since the players have nothing to say to each other.)
2. Suppose that in a symmetric game—a game in which the payoff matrix looks exactly the same from both players' point of view—player *I* selects a strategy first. Player *II* picks a strategy after seeing what player *I* has done. You would think that in a game where the players have identical roles except that player *II* has the benefit of some additional information, player *II*'s position would be at least as good as player *I*'s. In the zero-sum game, this situation could not

possibly be to player *I*'s advantage. In the non-zero-sum game, however, it may. Let us carry the point a little further. It is often to a player's advantage to play before an opponent, even if the rules do not require that he or she do so, or, alternatively, to announce the chosen strategy so that his or her decision becomes irrevocable.

3. Suppose that the rules of a game are modified so that one of the players can no longer pick some of the strategies that were originally available to him or her. In the zero-sum game, the player might not lose anything, but he or she certainly could not gain. In the non-zero-sum game, the player might very well gain.
4. It often happens in a non-zero-sum game that a player gains if an opponent does not know his or her utility function. This is not surprising. What is surprising is that at times it is an advantage to have your opponent know your utility function, and he or she may be worse off after learning it. This doesn't come up in zero-sum games: there it is assumed that each player knows the other's utility function.

Communication

The extent to which players can communicate has a profound effect on the outcome of a game. There is a wide spectrum of possibilities here. At one extreme we have games in which there is no communication whatever between the players and the game is played only once (later I will discuss why this matters). At the other extreme are the games in which the players can communicate freely. Generally, the more cooperative the game—the more the players' interests coincide—the more significant is the ability to communicate. In the zero-sum game that is completely competitive, communication plays no role at all. In the completely cooperative game, the problem is solely one of communication, and so the ability to communicate is crucial.

Cooperative games in which the players can communicate freely present no conceptual difficulty. There may be technical difficulties, of course, such as exist when a control tower directs a pilot in heavy traffic. But two players who can't communicate directly, such as two sailboat captains trying to avoid a collision in choppy waters or two guerrillas behind enemy lines, have a problem.

In completely cooperative games, communication is an unalloyed blessing. In games in which the players have some conflicting interests, communication has a more complex role. To see this, look at the game shown in figure 5.11. There are two basic properties of this game that are worth noting. No matter which strategy a partner plays, *a player always does best if he or she plays strategy* B. Also, no matter what strategy a player finally selects, *he or she will always do better if the partner plays strategy* A. Putting these two together, we have a conflict: on the one hand, each player should play strategy *B;* on the other hand, each is hoping the other will play strategy *A*. It is to the advantage of both, moreover, that the conflict be resolved by both choosing *A* rather than *B*. How do the players go about getting this result?

Figure 5.11

		PLAYER II: *A*	PLAYER II: *B*
PLAYER I	*A*	(5, 5)	(0, 6)
	B	(6, 0)	(1, 1)

If the game is played only once, the prospects for the (5, 5) outcome—sometimes called the cooperative outcome—are pretty bleak. Neither player can influence the partner's play, so the best each can do is play *B*. When the game is played repeatedly, it is a little different; it may be possible then to

play in such a way that one's partner is induced to play *A*. And eventually it may be possible to get (5, 5) payoff.

Granted there is an opportunity to arrive at this payoff when the game is played repeatedly, how do you invite your partner to cooperate in a specific situation? One way to do it is unilaterally to adopt an apparently inferior strategy (such as strategy *A* in this illustration) and hope the other player catches on. If your partner turns out to be either stupid or stubborn, there is nothing to do but revert to the cutthroat strategy, *B*, and settle for a payoff of 1. Thus, in fact, it may be possible to communicate in a game even if there is no direct contact between the players.

This kind of tacit communication is generally ineffective, however, most likely because it is generally misunderstood. Merrill M. Flood (1952) described an experiment in which a game, similar to the gasoline price war discussed earlier, was played a hundred times by two sophisticated players. The players made written notes during the game, and these left little doubt that the players were not communicating. This same conclusion was also reached in other experiments, as will be seen later.

Even if the message requesting cooperation does get through, it may not be in the best interests of the receiver to accept it immediately. It has been suggested that the appropriate response is to "misunderstand" at first, and let the other player "teach" you; then, before the other player loses hope, join him or her in playing the cooperative strategy. Following this strategy, in the example you would make 6 while "learning" and 5 thereafter. There is some evidence that this is what actually happened in the Flood experiment. But this is a dangerous game, as shall be seen later in the discussion of computer play in prisoner's dilemma games.

An amusing way to exploit the inability of players to get together was suggested by Luce and Raiffa (1957). They

would have a company offer each of two players two alternative strategies: a "safe" strategy and a "double-cross" one. If both play safe, they each receive a dollar; if both play double-cross, they each lose a nickel; and if one plays safe while the other plays double-cross, the safe player gets a dollar and the double-cross player gets a thousand dollars. The matrix in figure 5.12 shows the possibilities. Luce and Raiffa felt that as long as the players were prevented from communicating and only played the game once, the company would get some free advertising at no great expense.

Figure 5.12

		PLAYER II: SAFE	PLAYER II: DOUBLE-CROSS
PLAYER I	SAFE	($1, $1)	($1, $1,000)
	DOUBLE-CROSS	($1,000, $1)	(−5¢, −5¢)

On the basis of the games we have seen up to now, the ability to communicate would seem to be an advantage. So far, the only messages contemplated are offers to cooperate, and these obviously must be in the interest of both players. Otherwise the offers would not be made or, if made, would not be accepted. But it is also possible to communicate threats. The game shown in figure 5.13 was devised by Luce and Raiffa. Compare what happens when the players can communicate with what happens when they can't.

If the players can't communicate, they obviously can't threaten. Player *I* can do no better than to play strategy *A*,

Figure 5.13

		PLAYER II: *a*	PLAYER II: *b*
PLAYER I	*A*	(1, 2)	(3, 1)
	B	(0, −200)	(2, −300)

and player *II* can do no better than to play strategy *a*. But when the players are allowed to communicate, there is a radical change. Just what will happen is not altogether clear; to some extent, it will depend on circumstances that haven't been discussed. But, in any case, player *II*'s position deteriorates. If agreements made by the players can be enforced, player *I* can threaten to play strategy *B* unless player *II* commits herself to playing strategy *b*. If player *II* submits, player *I* will gain 2 and player *II* will lose 1, relative to the original (1, 2) payoff. If side payments are allowed, player *II*'s position becomes even worse. Not only can player *I* dictate player *II*'s strategy, but he can demand something under the table as well. True, player *II* can refuse to negotiate with player *I* and simply ignore his threats. It would still be in player *I*'s best interests to play strategy *A*, but who can say what would have a stronger influence on his behavior, his self-interest or his pique at being rebuffed? The point is that player *II* can avoid all these complications if communication is impossible.

The Order of Play

In the zero-sum game, the players pick their strategies simultaneously, neither player knowing the opponent's choice. If a player manages to find out the opponent's strategy in advance, he or she is far ahead and, in principle at least, the game is too trivial to be of interest. The non-zero-sum game is altogether different. The game may be far from trivial even when one player learns the opponent's strategy. What's more, the advantage of having this information may turn out to be a disadvantage. Let's look at an example.

A buyer and a seller are negotiating a contract in which the price per item and the quantity to be sold are still to be

determined. According to established procedure, the seller first sets the price, which, once quoted, may not afterward be changed; and the buyer indicates the quantity he wishes to buy.

In this instance, a wholesaler can buy two items from the manufacturer, one for $4 and the other for $5. The retailer has two customers for these items, one of whom is willing to pay $9 and the other, $10. If the mechanism for negotiating is as has just been described, what strategies should the players adopt? What will be outcome be?

Certain obvious characteristics of this game should be noted. It is clearly to both players' advantage to get together and in some way share the potential $10 profit. Also, if they are to share this profit equally and obtain a "fair" outcome, the selling price should be set at $7.

The wholesaler may have her eye on something better than a fair outcome, however. If she sets the price at $8 rather than $7, it would still be in the retailer's best interests to buy both items, although he would make a profit of only $3 rather than $5. (If he buys only one item, his profit would be $2, and if he buys nothing, he would have no profit at all.) In effect, the negotiating mechanism that requires the wholesaler to make the first move allows her to put pressure on the retailer, to her own advantage.

True, the retailer need not act mechanically in his own "self-interest" and allow himself to be exploited. Moreover, in the pure bargaining game in which buyer and seller negotiate freely about quantity and price simultaneously, the players will not always arrive at a selling price of $7; personality factors may affect the price one way or the other. Even though the consequences of requiring the wholesaler to move first cannot be predicted precisely, however, the general effect is, clearly, to give the wholesaler the upper hand.

The Effect of Imperfect Information

The wholesaler-retailer game just described was one of a series of experiments Lawrence E. Fouraker and Sidney Siegel (1963) studied in the laboratory. In the series, a variation of the game was also played. The original buying and selling prices were kept the same, but the rules were changed somewhat and, with them, the character of the game. In the new version, the wholesaler knew only her own profits; the retailer knew both players' profits. In addition, the retailer knew that the wholesaler did not know.

The basic difference between the variation and the original game was the reaction of the retailer when the wholesaler quoted a high selling price. In the original game, in which both players were fully informed, a high price on the wholesaler's part was interpreted by the retailer as a sign of greed, and as a result he often refused to go along. But when the retailer knew that the wholesaler was unaware of how large a share of the total profit she was asking, the retailer generally accepted his fate philosophically and did the best he could under the circumstances. Thus the wholesaler often did better—and the retailer, worse—when she had less information.

A corollary of this is that it is often to a player's advantage to see that the partner is well informed. Suppose, in a labor–management dispute, that labor's demands are such that if granted, they would force the company out of business. In such an instance, the company should see to it that the union is informed of the effect of its demands. Of course, the company's interest is not so much to have the union know the truth for its own sake as it is to have the union believe that its goals are unattainable; and for this purpose, a lie will do just as well. The company, then, may try to deceive the union about, say, its inability to compete if it grants a raise. Or it may lie about its utility function, stating that it would prefer a

prolonged strike to granting a raise, when in fact it would not. The union, for its part, might exaggerate the size of its strike fund. Ralph Cassady, Jr., (1957), describes some of the tactics used by competitors in a taxicab price war to "inform" (misinform) their opponents. They printed signs—which were never intended to be used—quoting rates considerably below the prevailing ones, and left them where they were sure to be seen by the competition. They also leaked the "information" that the company's owner had inherited a fortune—actually, he had inherited only a modest sum of money—and so gave the impression that the company had the capacity and the intention of fighting for some time.

In such a situation, players can gain if they can convince their opponent that they have certain attitudes or capabilities, whether they really have them or not. (If you are bargaining for an antique that you particularly want, it is a good idea not to let the seller know it.) If players really have these capabilities or attitudes, we have the situation spoken of earlier, where a player gains when the partner is better informed.

The Effect of Restricting Alternatives

Players are sometimes prevented from using some of their strategies. It is one of the paradoxes of non-zero-sum games that this restriction of a player's choice may be turned to that player's advantage. On the face of it, this seems absurd. How is it possible for a player to gain if prevented from using certain strategies? If it were to the player's advantage to avoid using certain strategies, could he or she get the same effect simply by not using them? No, the player could not; pretending you don't have certain alternatives is not the same as not having them. Suppose, for example, in the labor–

management dispute, that there were strict wartime controls in effect that prevented the company from raising wages. In that case, a union that under different circumstances might go on strike would very likely continue working, without objection. By the same token, it would be to the woman's advantage, in the "battle of the sexes" example, if she tended to faint at the sight of blood and could not, therefore, sit through a prize fight.

When circumstances do not restrict players' alternatives, they may try to gain an advantage by unilaterally restricting them themselves. This is not always so effective, however. If the woman in the "battle of the sexes" commits herself and her husband to the ballet by buying two tickets in advance, her husband may refuse to accompany her out of pique. If she cannot attend a fight because she faints at the sight of blood, a factor over which she has no control, he may view it differently.

This principle of limiting your alternatives to strengthen your position can be applied in other ways. We have already seen one application in the "business partnership" example. Another application is the hypothetical weapon called the doomsday machine. A weapon of great destructive power, it is set to go off *automatically* whenever the nation that designed it is attacked.

The point of building a doomsday machine is this: so long as the defending nation preserves its option to withhold retaliation, the way is open to an unpunished attack. A potential aggressor may be tempted to attack and then, by threatening an even worse attack, inhibit the defending nation from retaliating. With the doomsday machine, the defending nation forecloses one of its options; it cannot help but retaliate. Here again, it is to a player's advantage to keep the partner well informed; if you have a doomsday machine, it is a good idea to let everyone know it.

Threats

A *threat* is a statement that you will act in a certain way under certain conditions. It is like the doomsday machine in that it limits your future actions: "If you cut your price by five cents, I'll cut mine by a dime." But it is different because this self-imposed restriction isn't binding; you can always change your mind. The purpose of a threat is to change someone's behavior: to make that person do something he or she would not do otherwise. If the threat is carried out, it will presumably be to the detriment of the party that is threatened, but often it is also to the disadvantage of the party making the threat.

A threat is effective only to the extent that it is plausible. The greater the price the party making the threat must pay to carry it out, the less plausible the threat. This leads to the following paradox: if the penalty to the people making the threat is very high, they will be reluctant to burn their bridges behind them by committing themselves irrevocably to the threat. But it is precisely the threatener's failure to burn the bridges that is the greatest inducement to ignoring the threat. When you are bargaining for a new car, you can discount such statements as "I won't let it go for a penny under $2,000." If the seller is convinced that at this price no sale will be made, the seller may later choose to ignore his or her own threat. In a store where prices are fixed by the management, however, and the final decision is made by a disinterested salesperson on salary, the "threat" of not lowering prices is essentially irrevocable. Of course, in that case the store must accept the possibility that it will occasionally lose a sale. In the "battle of the sexes," one of the players may threaten to go it alone, but, again, the threat is not irrevocable, and the other player may threaten to resist. In the wholesaler-retailer example, the wholesaler was committed to her

quoted selling price by the rules of the game, and, as a result, the threat was much stronger.

Often both players are in a position to threaten. In a bargaining game, for example, both the buyer and the seller can refuse to complete the sale unless the price is right. In the battle of the sexes, each player could threaten to attend his or her preferred entertainment alone. Occasionally, however, it happens that only one of the players is in a position to threaten. Such a situation was cleverly exploited by Michael Maschler (1963) in his analysis of an inspection model designed to detect illicit nuclear testing.

The two players represented in this model are countries that have signed an agreement outlawing nuclear testing. One of them is considering violating the treaty; the other wants to be able to detect a violation if it occurs. (In reality, a country may play either role, or both at the same time.) The inspecting country has the benefit of detecting devices that indicate natural as well as artificial disturbances; it is allotted a quota of on-site inspections. The strictly mathematical problem consists of timing the inspections so as to maximize the probability of discovering a violation if there is one and determining when, if ever, the potential treaty violator would test a nuclear device.

This, obviously, is not a zero-sum game, since presumably both the inspector and the inspected would prefer that there be no violations to having violations that are subsequently discovered. In any case, this is assumed in the model. Maschler showed that the inspector actually does best by announcing a strategy in advance and keeping to it, much as the wholesaler gained by setting a high price. (This is based on the assumption that the violator believes the announcement and then acts according to his or her own self-interest. There is no reason why the violator shouldn't believe it, for it is in the inspector's self-interest to tell the truth.) Why can't

the tester use similar tactics; that is, announce an intention to cheat in some set manner and have the inspector make the best of it? If you look at the bare payoff matrix, it is clear that the tester can, but the political realities are such that this is not feasible.

Henry Hamburger (1979) tells how a community might use this same principle to enforce its speeding laws. He estimated the utility to the potential speeder and community of each of a number of factors: (1) the time saved by speeding; (2) the risk to the driver; (3) the driver's fine when caught; (4) the cost of enforcement, and (5) the hazard to others of speeding—and arrives at the matrix shown in figure 5.14.

Figure 5.14

		COMMUNITY	
		ENFORCE	IGNORE
DRIVER	VIOLATE	(−190, −25)	(10, −5)
	COMPLY	(0, −20)	(0, 0)

It appears that whatever the driver does, the community does better to ignore the speeding; in such a case the driver can be expected to ignore the laws and the community's payoff will be minus 5. But if the community announces and implements a policy of enforcing the law 10 percent of the time, the situation can be depicted in figure 5.15.

Figure 5.15

		COMMUNITY
DRIVER	VIOLATE	(−10, −7)
	COMPLY	(0, −2)

Now it is to the driver's advantage to comply and the community's loss is only 2.

Once again, looking only at figure 5.14, the driver could

also adopt a strategy in advance to force the community's strategy; however, in the real world this is not a practical alternative.

Binding Agreements and Side Payments

When players negotiate, they often reach some sort of agreement. In some games there is no mechanism for enforcing agreements and the players can break their word with impunity. But in other games this is not the case. Though any agreement that is reached is reached voluntarily, once made, it is enforced by the rules. This possibility of making *binding agreements* has a strong influence on the character of the game.

Let us look at an anecdote related by Merrill M. Flood (1952). Though this "game" actually had more than two players, it is pertinent here.

Flood wanted one of his children to baby-sit the others. He suggested that a fair way of selecting the sitter and setting a price for the service was to have the older children bid against one another in a backward auction. That is, he would start with a price of $4—the most he was willing to pay—and the children would bid successively, each bid lower than the last one, until the bidding stopped. The last person to bid would baby-sit for the agreed price.

It wasn't long before the children hit upon the possibility of collusion. When they asked about it, their father said he would allow them to "rig" the bidding if they satisfied two conditions: the final price must be no more than the ceiling of $4 he had set originally, and the children must agree among themselves in advance who would do the baby-sitting and how the money would be divided. As it happened, the children did not reach an agreement. Several days later, a bona fide auction was held and the final price was set at 90 cents. Thus the mere opportunity to communicate and the existence of a mechanism to enforce agreements are not sufficient to guarantee that an agreement will be reached. In this game the failure to get together resulted in an outcome markedly inferior to what might have been if the players had acted together.

In some games it is possible for one player to affect the actions of another by offering a "side payment," a payment made "under the table." This was the case in Flood's baby-sitting example. If an agreement had been reached, the child who baby-sat would have paid the other children a certain amount in return for staying out of the bidding.

In many games, however, the players cannot or may not make side payments. Sometimes this is a matter of policy—when the government invites private companies to bid on a contract, it responds to a cooperative strategy on the bidders' part (and to the corresponding side payments) with less than enthusiasm. Sometimes side payments are impractical because there is no unit that can be transferred from one player to the other. In the "battle of the sexes," the pleasure the wife feels when her husband takes her to the ballet is simply not transferable. (She may be able to reciprocate in a future game, however; perhaps next time she'll go to the fights.) Similarly, one legislator is barred from giving another any direct payment such as money in return for support, but the legislator may repay the supporter in kind in the future.

The simple example in figure 5.16 will help clarify the role played by side payments. If it is impossible to make side payments, player *II* can do no better than choose strategy *B* and take a dollar. But if side payments are possible (and if the players can make binding agreements), it is a very different game. Player *II* is then in a position to demand a sizable

portion of player *I*'s thousand dollars, and if player *I* refuses, he may be left with only $100. Whether player *II* will actually follow through on her threat and sacrifice a dollar, player *I* must decide for himself.

Figure 5.16

		PLAYER II: *A*	PLAYER II: *B*
PLAYER I	*a*	($100, 0)	($1,000, $1)

You cannot predict with any degree of assurance what will happen in this kind of game, or even to prescribe what should happen in theory. In practice, what happens no doubt depends strongly on the prize at stake. In games 1 and 2 in figure 5.17, the players may communicate and may make binding agreements, but side payments are not permitted.

Figure 5.17

GAME 1

		PLAYER II: *A*	PLAYER II: *B*
PLAYER I	*a*	(0, 0)	($1,000, $10)
	b	(0, 0)	($900, $1,000)

GAME 2

		PLAYER II: *C*	PLAYER II: *D*
PLAYER I	*c*	(0, 0)	($1,000, $800)
	d	(0, 0)	($900, $850)

In game 1, player *II* will very likely try to get $1,000 by playing strategy *B* and persuading player *I* to play strategy *b*. Player *II*'s threat is to play strategy *A*, which will reduce player *I*'s payoff to nothing, at a cost of only $10 to himself.

Prudence seemingly dictates that player *I* should settle for the $900.

In game 2, the discussion might go much the same way, but now player *II*'s threat is less plausible. Her "spite" strategy, strategy *C*, is a lot more expensive, though her threat to player *I* is essentially the same. Still, it is difficult to say what the players will do. If player *I* does decide to call player *II*'s bluff, however, he plays a more dangerous game in 1 than in 2.

An Application to a Prison Break

A vivid illustration of some of these factors at work in real life was detailed in a news article in the *New York Times* on June 28, 1965. The article described a prison disturbance in which two guards were taken hostage. The warden refused to negotiate with the prisoners as long as the guards were held captive, and the guards were eventually freed, unharmed. The warden was quoted as saying: "They wanted to make a deal. I refuse to make deals at any time. Therefore I didn't listen, and I don't know what the deal was about." Figure 5.18 analyzes the "game."

Figure 5.18

		THE WARDEN: NEGOTIATE	THE WARDEN: DO NOT NEGOTIATE
PRISONERS	HARM GUARDS	*A*	*C*
	DO NOT HARM GUARDS	*B*	*D*

We can start by ruling out *A*, since the prisoners gain nothing by harming the guards if they are released. Of the three remaining possibilities, the prisoners most prefer *B*,

then *D*, and least of all, *C*, which would mean additional punishment without any compensating gain. The warden wanted *D* the most, then *B* (presumably), and finally *C*.

The only chance of negotiating the prisoners had was to threaten to harm the guards unless they were released and to hope that the threat would be believed. But the warden simply cut communications; he refused to hear "what the deal was about." In effect, he committed himself to a strategy, not freeing the prisoners, and forced the prisoners to make a choice between *C* and *D*. His hope was that they would make the best of a bad situation by choosing *D* rather than *C*. And so they did—but if they had been sufficiently vindictive, they might have acted otherwise. Whether the warden made the proper choice is a question about which reasonable people might differ. Indeed, the warden and the hostages probably did.

The Prisoner's Dilemma

Two men suspected of committing a crime together are arrested and placed in separate cells by the police. Each suspect may either confess or remain silent, and each one knows the possible consequences of his action. These are: (1) If one suspect confesses and his partner does not, the one who confessed turns state's evidence and goes free and the other one goes to jail for twenty years. (2) If both suspects confess, they both go to jail for five years. (3) If both suspects remain silent, they both go to jail for a year for carrying concealed weapons—a lesser charge. We will suppose that there is no "honor among thieves" and each suspect's sole concern is his own self-interest. Under these conditions, what should the criminals do? The game, shown in figure 5.19, is the celebrated

prisoner's dilemma which was originally formulated by A. W. Tucker and which has become one of the classical problems in the short history of game theory.

Figure 5.19

		SUSPECT II: CONFESS	SUSPECT II: DO NOT CONFESS
SUSPECT I: CONFESS		(5 yrs., 5 yrs.)	(go free, 20 yrs.)
SUSPECT I: DO NOT CONFESS		(20 yrs., go free)	(1 yr., 1 yr.)

Let us look at the prisoner's dilemma from the point of view of one of the suspects. Since he must make his decision without knowing what his partner will do, he must consider each of his partner's alternatives and anticipate the effect of each of them on himself.

Suppose his partner confesses; our man must either remain silent and go to jail for twenty years, or confess and go to jail for five. Or, if his partner remains silent, he can serve a year by being silent also, or win his freedom by confessing. Seemingly, in either case, *he is better off confessing!* What, then, is the problem?

The paradox lies in this. Two naïve prisoners, too ignorant to follow this compelling argument, are both silent and go to prison for only a year. Two sophisticated prisoners, primed with the very best game-theory advice, confess and are given five years in prison in which to contemplate their cleverness.

We will return to this argument in a moment, but before we do, let's consider the essential elements that characterize this game. Each player has two basic choices: he can act "cooperatively" or "uncooperatively." When all the players act cooperatively, each does better than when all of them act uncooperatively. For any fixed strategy (ies) of the other player(s), a player always does better by playing uncooperatively than by playing cooperatively.

In the following examples, taken from many different contexts, these same basic elements appear.

1. Two different firms sell the same product in a certain market. Neither the product's selling price nor the total combined sales of both companies vary from year to year. What does vary is the portion of the market that each firm captures, and this depends on the size of their respective advertising budgets. For the sake of simplicity, suppose each firm has only two choices: spending $6 million or $10 million. The size of the advertising budget determines the share of the market and, ultimately, the profits of each company, as follows:

If both companies spend $6 million, they each get a $5 million profit. If a company spends $10 million when its competitor spends only $6 million, its profit goes up to $8 million at the expense of its competitor, who now loses $2 million. And if both companies spend $10 million, the extra marketing effort is wasted, since the market is fixed and the relative market position of each company remains the same; consequently, the profit of each company drops to $1 million. No collusion is allowed between the firms. The game is shown in figure 5.20.

Figure 5.20

		CORPORATION II	
		SPEND $6 MILLION	SPEND $10 MILLION
CORPORATION I	SPEND $6 MILLION	($5 million, $5 million)	(−$2 million, $8 million)
	SPEND $10 MILLION	($8 million, −$2 million)	($1 million, $1 million)

2. There is a water shortage and citizens are urged to cut down on water consumption. If each citizen responds to these requests by considering his or her own self-interest, no

one will conserve water. Obviously, any saving by an individual has only a negligible effect on the city's water supply, yet the inconvenience involved is very real. On the other hand, if everyone acts in his or her own self-interest, the results will be catastrophic for everyone.

3. If no one paid taxes, the machinery of government would break down. Presumably all citizens would prefer that everyone, including themselves, pay taxes to having no one pay taxes. Better yet, of course, everyone would pay taxes except the individual him- or herself.

4. After several years of overproduction, farmers agree to limit their output voluntarily in order to keep the prices up. But no one farmer produces enough seriously to affect the price, so each starts producing what he or she can and selling it for what it will bring, and once again there is overproduction.

5. Two unfriendly nations are preparing their military budgets. Each nation wants to obtain a military advantage over the other by building a more powerful army, and each spends accordingly. They wind up having the same relative strength, and a good deal poorer.

As we can see from these examples, this kind of problem comes up all the time. For convenience's sake, let us fix our attention on a single game, using the first example, that of two companies setting their advertising budgets. The game will be the same, except for an additional condition. The budget will not be set once but, more realistically, we will assume that it is fixed annually for a certain period of time; say, twenty years. When each company decides on its budgets for any given year, it knows what its competitor spent in the past.

In discussing the prisoner's dilemma, we had decided that if the game is played only once, the prisoners have no choice but to confess. The same line of reasoning applied here leads

to the conclusion that the firms should each spend $10 million. But when the game is played repeatedly, the argument loses some of its force. It is still true that if you spend $10 million in a given year, you will always do better than if you spend $6 million *that year*. But if you spend $10 million in one year, you are very likely to induce your competitor to spend $10 million the next year, and that is something you don't want. A more optimistic strategy is to signal your intent to cooperate by spending $6 million and hope your competitor draws the proper inference and does the same. This strategy could lead to a cooperative outcome and in practice often does. But in theory there is a problem.

The argument that spending $6 million in any one year tends to encourage your competitor to spend the same amount the next year is all very well for the first nineteen years, but it clearly breaks down in the twentieth year. In the twentieth year, there is *no next year*. When the firms reach the twentieth year, they are in effect in the same position they are in when they play the game only once. If the firms want to maximize their profits, and we assume that they do, the argument favoring the uncooperative strategy is as compelling now as it was then.

But the argument doesn't end there. Once the futility of cooperating in the twentieth year is recognized, it follows that there is no point in cooperating in the nineteenth year either. And if there is no chance of inducing a cooperative response in the nineteenth year, why cooperate in the eighteenth year? Once you fall into the trap, you fall all the way: there is no point in cooperating in the eighteenth, the seventeenth . . . or the first year either. If you accept the argument supporting the uncooperative strategy in the single instance, it follows that you must play uncooperatively not only in the last of a series but in every individual trial as well.

It is when the prisoner's dilemma is played repeatedly—

and not for a fixed number of trials but for an indefinite period—that the cooperative strategy comes into its own. And these are precisely the conditions under which the prisoner's dilemma is often played. Two competing firms know that they won't be in business forever, but they generally have no way of knowing when death, merger, bankruptcy, or some other force will end their competition. Thus the players can't analyze what will happen in the last "play" and then work backward from there, for nobody knows when the last "play" will be. The compelling argument in favor of the uncooperative strategy breaks down, then, and we breathe a sigh of relief.

This is really the point. The prisoner's dilemma has one characteristic that makes it different from the other games that have been discussed. As a rule, when analyzing a game, you are content if you can say what rational players should do and predict what the outcome will be. But in the prisoner's dilemma the uncooperative strategy is so unpalatable that the question most people try to answer is not What strategy should a rational person choose? but How can we justify playing a cooperative strategy? Many different answers to this last question have been proposed. Let us consider a few.

The Prisoner's Dilemma in the Past

The prisoner's dilemma is important because it presents in capsule form a problem that arises in many different ways. Some of the manifestations of the prisoner's dilemma were discussed long before there was a theory of games. Thomas Hobbes, the political philosopher, examined one version of the dilemma in which the "players" were the members of society.

Society, Hobbes conjectured, was originally in a state of anarchy. Constant warfare and banditry were the consequence of each individual's trying to further his own narrow self-interest; it was a society in which *B* might murder *C* for a shiny bauble, and *B* might be murdered in turn for the same reason. It would be to everyone's individual advantage, Hobbes felt, if restrictions were imposed and enforced; that is, *B* would prefer to give up the chance of getting the bauble in return for additional security. Hobbes saw the social contract as an enforced cooperative outcome. In *The Leviathan* (1962) he described the creation of a government (preferably a monarchy) "as if every man should say to every man, I authorize and give up my right of governing myself, to this man or to this assembly of men, on this condition that thou give up thy right to him, and authorize all his actions in a like manner" (p. 104). (Whether historically this is an accurate picture is unimportant; what matters is how the problem was perceived and how it was resolved.)

Hobbes, after pointing out the disadvantages of the uncooperative outcome, suggests that the many independent decisions to cooperate or not be taken out of the hands of the people that make up society. In effect, society should submit to compulsory arbitration, and the government should play the role of the arbiter. This is not an uncommon point of view. Luce and Raiffa, in *Games and Decisions* (1957), make the same point: "Some hold the view that one essential rule of government is to declare that the rules of social 'games' be changed whenever it is inherent in the game situation that the players, in pursuing their own ends, will be forced into a socially undesirable position" (p. 97).

On a more modest scale, the noted sociologist George Simmel (1955) recognized that competing businesses are often faced with what amounts to a prisoner's dilemma. He

described the behavior of businessmen who, in effect, are playing this game for an indefinite number of times:

> Inter-individual restriction of competitive means occurs when a number of competitors voluntarily agree to renounce certain practices of outdoing one another—whereby the renunciation by one of them is valid only so long as the other too observes it. Examples are the arrangement among booksellers of a given community to extend none or more than five or ten percent discount; or the agreement among store owners to close their businesses at eight or nine o'clock, etc. It is evident that here mere egoistic utility is decisive: the one does without those means of gaining customers because he knows that the other would at once imitate him if he did not, and that the greater gain they would have to share would not add up to the greater expenses they would likewise have to share. . . . In Economics, the third party is the consumer; and thus it is clear how the road toward cartelization is taken. Once it is understood that one can do without many competitive practices provided the competitor does likewise, the result may not only be an even more intense and purer competition, which has already been emphasized, but also the opposite. The agreement may be pushed to the point of abolishing competition itself and of organizing enterprises which no longer fight for the market but supply it according to a common plan. . . . This teleology, as it were, transcends the parties, allowing each of them to find its advantage and achieves the seeming paradox that each of them makes the opponent's advantage its own. (p. 76)

Note that Simmel looks at the consumer as an outsider, who, though affected by what the competing businesses do, has no control over what happens. In effect, the consumer is not a player. While cooperation between businesses is easily turned to their mutual advantage, the effects can easily be antisocial as far as society as a whole is concerned. And so society prohibits "cooperative play" in the form of trusts, cartels, price fixing, and bribery.

This theme—the tendency of competitors to avoid mutually destructive price competition—is also considered by John Kenneth Galbraith. In *American Capitalism: The Concept of Countervailing Power* (1952), he says: "The convention against price competition is inevitable. . . . The alternative is self-destruction" (p. 112). Price competition, then, tends to be replaced by competition in sales and advertising. When there are only a few sellers, interindustrial competition diminishes. If a few large industries are bargaining with their workers on wages, there ceases to be a wage competition among the individual industries to attract workers. There is, rather, bargaining in between the "countervailing powers" of labor and management.

Problems of the prisoner's dilemma type, in one form or another, have been around for some time. The cooperative strategy is generally accepted as the "proper" one (except when its effect is antisocial), sometimes for ethical reasons. Immanuel Kant (1959) asserted that a person should decide if an act is moral by examining the effect of everyone's acting similarly; the golden rule says much the same thing. More recently Anatol Rapoport in *Fights, Games and Debates* (1960) asserted that there are considerations other than a player's narrow self-interest that he or she should take into account when choosing play. If there is to be any hope of reaching the elusive cooperative outcome, Rapoport feels, it is necessary that the players accept certain social values, and having accepted these values, players should cooperate even in the one-trial, version of the prisoner's dilemma. His argument goes like this:

> Each player presumably examines the whole payoff matrix. The first question he asks is "When are we both best off?" The answer in our case is unique: at [the cooperative outcome]. Next "What is necessary to come to this choice?" Answer: the assumption that whatever I do, the other will, made by both parties. The conclusion is, "I am one of the parties; therefore I will make this assumption." (p. 177)

Rapoport is well aware that his point of view is in conflict with "'rational' strategic principles" of self-interest; he simply rejects these principles. The minimax strategy in the two-person, zero-sum game, he asserts, is also based on an assumption: that one's opponent will act rationally; that is, in accordance with his or her own self-interest. If in this game one's opponent fails to act rationally, the minimax strategy will fail to exploit the opponent's errors. Just as you may be mistaken when you assume your opponent will be rational in the zero-sum game, you may be mistaken in the non-zero-sum game when you anticipate your partner's goodwill.

Since most people are generally reluctant to accept the uncooperative strategy as the proper one, there is a temptation, whenever a way out is proposed, to take it. Despite my sympathy with Rapoport's effort, however, I do not believe the paradox of the prisoner's dilemma has really been overcome. To see why, let us go back a bit.

In the prisoner's dilemma examples discussed in this chapter, the payoffs were stated in terms of "years in prison" or net profits, rather than utiles. The actual utilities involved are only suggested by such statements as: "Each player is only concerned with his own self-interest," or "Each company only wants to maximize its own profits." For the sake of simplicity, a more formal description in terms of utilities was avoided. But at bottom are certain assumptions that, though stated only inexactly, are critical. If they are not valid, we may be playing a radically different game from the one we think we are playing. If, in the original prisoner's dilemma, a prisoner would prefer to spend a year in jail along with his partner rather than go free, knowing his partner was serving twenty years, the argument for confessing breaks down. But

then the game could hardly be called a prisoner's dilemma. The original paradox has in fact not been resolved. This is the objection to Rapoport's argument: his golden rule approach assumes the problem out of existence. If players are as concerned with a partner's payoff as they are with their own, the game isn't a prisoner's dilemma; and if each player is interested solely in his or her own payoff, Rapoport's comments aren't pertinent.

The analogy between the assumptions that are made in the zero-sum and in the non-zero-sum games is not very convincing either. In the zero-sum game, you can get the value of the game whether your opponent is good, bad, or indifferent; you do not have to assume that he or she is rational. Rapoport states as much but goes on to say that by playing minimax, you lose the opportunity of exploiting errors, and with a foolish opponent you should not be content to get only the value of the game.

This is not entirely true. There are games—we have seen some of them already—in which an inferior strategy on one player's part and the minimax strategy on the other's will lead to a payoff that is greater than the value of the game for the minimax player. But even in those games in which playing minimax precludes your getting more than the value of the game, the analogy is questionable. In order to exploit your opponent's weakness, it is not enough to know he or she will deviate from the minimax; you must also know how. Suppose, for example, you are going to match pennies *once* with a simpleton and you decide to assume that she will *not* play rationally. Specifically, you believe that she is quite capable of playing one side of the coin with a probability greater than one-half. But which side? And how can you exploit her inferior strategy if you don't know which side?

As a rule, when you play a minimax strategy in a zero-sum game it is not because you have faith in your opponent's

rationality but because you have no other, more attractive alternative—and this is true even when you suspect your opponent is capable of making a mistake.

In the prisoner's dilemma, however, the assumption that your partner will cooperate is really an assumption. Unless you are a masochist, if you play cooperatively you must believe as an act of faith that your partner will too. Even if your partner cooperates and seemingly justifies your act of faith, some players will still question your choice since you could have done better still by playing uncooperatively. This attitude may seem piggish, but, then, players don't look to game theorists for moral principles; they already have their own. All they ask is to find a strategy that will suit their purpose, selfish or otherwise.

The difference between the assumptions made in the zero-sum and in the non-zero-sum games is even more clear when they fail. And there is no question that people often *do* fail to play cooperatively in society's prisoner's dilemmas. In non-zero-sum games, cooperating with a partner who doesn't cooperate with you leads to disaster; in zero-sum games, the worst that can happen when you play minimax is that you lose an opportunity to swindle your opponent.

The Nash Arbitration Scheme

Players in a bargaining game are in an awkward position. They want to make the most favorable agreement that they can, while avoiding the risk of making no agreement at all; and, to a certain extent, these goals are contradictory. If one party indicates a willingness to settle for any terms, even if the gain is only marginal, he or she will very likely arrive at an agreement, but not a very attractive one. On the other

hand, if he takes a hard position and sticks to it, he is likely to reach a favorable agreement if he reaches any agreement at all—but he stands a good chance of being left out in the cold. A car dealer who is eager to sell a car will hide this fact from a buyer; but she will go to great lengths to determine how much she must drop her price to make a sale—even to the extent of using hidden microphones.

Even when a player is reconciled to receiving only a modest gain and pushes hard for an agreement, his or her eagerness is often interpreted as weakness, leading the partner to stiffen his or her demands and actually lessening the chance of an agreement. This is what often happens when one of two warring nations sues for peace. In a labor–management dispute that actually took place, one party who had a weak heart offered favorable terms in order to reach an immediate settlement. What happened was very different from what had been anticipated. Instead of waiving the usual bargaining procedure, the other party became suspicious and then resistant. The terms of the ultimate settlement were identical to those originally offered, but they were arrived at only after very hard bargaining.

One way of circumventing the actual bargaining process, at least in principle, is to have the terms of the agreement set by arbitration. In this way you can avoid the danger of not reaching a settlement at all. The problem is to establish arbitration that somehow reflects the strengths of the players in a realistic way, so that you get the effects of negotiation without risk. John Nash (1950) suggests the following procedure.

He begins by assuming that two parties are in the process of negotiating a contract. They might be management and labor, two countries formulating a trade agreement, a buyer and seller, and so forth. For convenience, and with no loss of generality, he assumes that a failure to agree—no trade, no sale, a strike, and so on—would have a utility of zero to both

players. Nash then selects a single arbitrated outcome from all the agreements that the players have it in their power to make: that outcome in which the product of the players' utilities is maximized.* This scheme has four desirable properties that he feels justify its use, and it is the *only* one that does. The four properties are:

1. *The arbitrated outcome should be independent of the utility function.* Any arbitrated outcome should clearly depend on the preferences of the players, and these preferences are expressed by a utility function. But, as we saw earlier, there are many utility functions to choose from. Since the choice of utility function is completely arbitrary, it is reasonable to demand that the arbitrated outcome not depend on the utility function selected.
2. *The arbitrated outcome should be Pareto optimal.* Nash considered it desirable that the arbitrated outcome be Pareto optimal; that is, that there should not be any other outcome in which both players simultaneously do better.
3. *The arbitrated outcome should be independent of irrelevant alternatives.* Suppose there are two games A and B in which every outcome of A is also an outcome of B. If the arbitrated outcome of B turns out also to be an outcome of A, this outcome must also be the arbitrated outcome of A. Put another way, the arbitrated outcome in a game remains the arbitrated outcome even when other outcomes are eliminated as possible agreements.
4. *In a symmetric game, the arbitrated outcome has the same utility for both players.* Suppose the players in the bargaining game have symmetric roles. That is, if there is an outcome that has a utility of x for one player and a utility of y for the other, there must also exist an outcome that has a

* It should be noted that players may not only obtain utility pairs associated with simple agreements but may also obtain intermediate pairs by using coordinated strategies as well. Suppose, for example, that in the battle of the sexes, going to the ballet has a utility of 4 for the husband and 8 for the wife, and going to the fights has a utility of 6 for the husband and 2 for the wife. Each may obtain a utility of 5 if they let a toss of a coin determine the choice of the evening's entertainment.

utility of y for the first player and x for the second player. In such a game, the arbitrated outcome should have the same utility for both players.

Before the Nash arbitration scheme can be applied, the utility function of both players must be known. This is its biggest disadvantage, for not only are the utilities not always known, they are often deliberately obscured by the players. If a player's utility function is misrepresented, it can be turned to his or her advantage. This is reassuring in a way, for it suggests the scheme is realistic. In real life also the utility function is often misrepresented, as we have seen.

It is important to realize that the Nash scheme is neither enforceable nor a prediction of what will happen. It is, rather, an a priori agreement obtained by abstracting away many relevant factors such as the bargaining strengths, of the players, cultural norms, and so on. (In this respect it is similar to the Shapley value that will be discussed later.) As a matter of fact, the Nash outcome often appears to be unfair: it tends to make the poor poorer and the rich richer. This is to be expected, however. A rich player is often in a stronger position than a poor one. To see how this works, consider the following example.

Suppose a rich woman and a poor woman can get a million dollars if they can agree on how to share it between them; if they fail to agree, they get nothing. In such a case, the Nash arbitration scheme would generally give the rich woman a larger portion than it would give the poor woman, because of the difference in their utility functions. Let's take a moment to see why.

When *relatively* large amounts of money are involved—that is, amounts of money that are large relative to what a person already possesses—people tend to play it safe. Most people, unless they are very rich, would prefer a sure million

dollars to an even chance of getting $10 million, although they would prefer an even chance of getting $10 to a sure dollar. But a large insurance company would prefer the even chance of $10 million, and in fact gladly accepts much less attractive risks every day. This indifference to the difference between large sums of money is reflected in the poor woman's utility function much more strongly than in the rich woman's. The relative attractiveness of $1 and $10 to the poor woman would be like the relative attractiveness of $1 million and $10 million to a woman who is very wealthy. A utility function that correctly reflects the poor woman's situation would be the square-root function: $100 would be ten utiles, $1 would be one utile, $16 would be four utiles, and so on. Thus the poor woman would be indifferent to an even chance at $10,000 and a sure $2,500. (The specific choice of the square-root function is of course arbitrary; many others would do as well.) It may be assumed that the rich woman's utility function is identical to the money in dollars. Under these conditions, Nash's suggested outcome would be that the rich woman gets two-thirds of the million dollars and the poor woman only a third.

Two-Person, Non-Zero-Sum Game Experiments

One reason for studying experimental games is that they are interesting. When you spend a lot of time thinking about how people should behave in theory, you become curious about how they actually behave in practice. A second reason for studying experimental games is that you may gain insights that will enable you to play better. That is a much more important consideration in non-zero-sum games than it is in zero-sum games. In the two-person, zero-sum game, players

can obtain the value of the game by their own efforts alone; they don't have to worry about what their opponent does. In the non-zero-sum game, unless you are willing to be satisfied with a minimal return—as a buyer and a seller may have to be when they fail to reach an agreement—you *must* be concerned with how your partner plays. Similarly, in a sequence of prisoner's dilemmas, what you anticipate will be your partner's play will affect your own.

Granted that it is worthwhile to learn more about how people behave, why should we turn to the laboratory rather than to actual life for our data? Examples of non-zero-sum games in everyday life are certainly common enough. Lawrence E. Fouraker and Sidney Siegel, in *Bargaining and Group Decision Making* (1960), answer the question thus:

> In the specific case of Bilateral Monopoly, it would be extremely unlikely that appropriate naturalistic data could be collected to test the theoretical models. This is not because the phenomenon is unusually rare. Indeed, there are numerous daily exchanges that are conducted under conditions that approximate Bilateral Monopoly: a franchised dealer negotiates with a manufacturer regarding quotas and wholesale price; two public utilities bargain about the division of some price they have placed on a joint service; a chain grocery store negotiates with a canner, who in turn must deal with farmers' cooperatives; labor leaders in a unionized industry deal with management in that industry; and so forth.

The trouble with "real" games is that they are not set up for our convenience. The variables are not controlled; we are not likely to find two situations that are identical except for one variable. This makes it difficult to determine just how important a variable is in influencing the final outcome. Also, it is usually not feasible to determine the payoffs. In the laboratory, on the other hand, the players can be separated to avoid personal contact (an unnecessary complicating factor),

the payoffs are clear, and the variables can be altered at will. It is also possible to motivate the players by making the payoffs sufficiently large—in principle, at least.

Some Experiments on the Prisoner's Dilemma

The many experiments conducted on the prisoner's dilemma all share a common purpose: to determine under what conditions players cooperate. Among the significant variables that determine how a player will behave are the size of the payoffs, the way the other person plays, the ability to communicate, and the personality of the players. In a series of experiments by Alvin Scodel, J. Sayer Minas, David Marlowe, Harvey Rawson, Philburn Ratoosh, and Milton Lipetz that were described in three articles in the *Journal of Conflict Resolution* between 1959 and 1962, a prisoner's dilemma game and some variations on it were played repeatedly. Let us consider some of the observations of the experimenters.

The game shown in figure 5.21, which I will call game 1, was played fifty times by each of twenty-two pairs of players. *C* and *NC* are the cooperative and noncooperative strategies, respectively. The players were physically separated throughout the fifty trials, so no direct communication was possible. At every trial, each player knew what his or her partner had done on every previous trial.

At each trial, each player had two choices, making a total

Figure 5.21

GAME 1

	C	*NC*
C	(3, 3)	(0, 5)
NC	(5, 0)	(1, 1)

of four possible outcomes. If the players had picked their strategies at random, we would have expected the cooperative payoff (3, 3) 25 percent of the time; the uncooperative payoff (1, 1) 25 percent of the time; and one of the mixed payoffs (5, 0) or (0, 5) 50 percent of the time. In fact, the uncooperative outcome predominated. Of the twenty-two pairs, twenty had more uncooperative payoffs than they had any other combination. Even more surprising, the tendency of the players was to be more uncooperative as the game progressed.

Game 1a was a repetition of game 1, with one variation: the players were allowed to communicate in the last twenty-five of the fifty trials. As you would expect, the results on the first twenty-five trials were pretty much the same as before. On the last twenty-five trials, there was still a tendency not to cooperate, but it was not so pronounced as when the players couldn't communicate.

Game 2 had the same payoff matrix as games 1 and 1a, but a variation was introduced. The subjects didn't play against each other but against the experimenter, though they weren't aware of it. The experimenter played in accordance with a predetermined formula: at each trial, he did what the subject did. If a subject cooperated, the experimenter did also (on the same trial), and the subject received 3. If the subject didn't cooperate, he received only 1. The game was repeated fifty times, just as games 1 and 1a were. The players chose the uncooperative strategy 60 percent of the time and played noncooperatively in the second twenty-five trials more often than they did in the first twenty-five trials.

You would think that in fifty trials the subjects would realize that they were not playing against just another person. And, presumably, subjects who caught on would play strategically. Both postexperimental interviews and the experimental results indicated, however, that every subject believed the

responses of his or her "partner" to be legitimate. Those players who noticed a similarity between their own and their "partner's" play attributed it to coincidence.

In several other prisoner's dilemma games, the same pattern was repeated. In game 3 (see figure 5.22), for example, a player only gained 2 by defecting from the cooperative outcome. Nevertheless, 50 percent of the time, in the first half of a thirty-trial sequence, the players were uncooperative. And in the last fifteen trials, the percentage went up to 65 percent. When the experiment was repeated with the second player replaced by the experimenter (who always played uncooperatively), the frequency of cooperative play was virtually unchanged. In game 4 (see figure 5.23), a prisoner's dilemma in which the payoffs were mostly negative and the players had to scramble to cut their losses, there was virtually no cooperation at all.

Figure 5.22

GAME 3

	C	*NC*
C	(8, 8)	(1, 10)
NC	(10, 1)	(2, 2)

Figure 5.23

GAME 4

	C	*NC*
C	(−1, −1)	(−5, 0)
NC	(0, −5)	(−3, −3)

Some of the most interesting experiments were not really prisoner's dilemma games at all. Three of these games are shown in figure 5.24. Each of the three games was played in thirty-trial sequences, and in each game it was surprising how frequently the players failed to cooperate. In game 5, the players failed to cooperate 6.38 times on the first fifteen trials

Figure 5.24

GAME 5

	C	NC
C	(6, 6)	(4, 7)
NC	(7, 4)	(−3, −3)

GAME 6

	C	NC
C	(3, 3)	(1, 3)
NC	(3, 1)	(0, 0)

GAME 7

	C	NC
C	(4, 4)	(1, 3)
NC	(3, 1)	(0, 0)

(on average) and 7.62 times on the last fifteen trials: a small but statistically significant increase.

Games 6 and 7 turned out about the same. In game 6, the players were uncooperative slightly more than half the time and were a little more cooperative during the first half of the trials than they were in the second half. In game 7, the players cooperated about 53 percent of the time, but in the last fifteen trials *they failed to cooperate more than half the time.*

Throughout these experiments, there was a marked consistent tendency to play uncooperatively. Uncooperative play is understandable in games of the prisoner's dilemma type, where it has obvious advantages, at least in the short run. But the tendency persisted into the last three games, and this is much more difficult to explain.

In game 5, uncooperative play was rewarded only part of the time: a player gained only if the partner cooperated—and not very much at that. And if the partner played uncooperatively also, the player received the smallest possible payoff. In games 6 and 7, it was absurd to play uncooperatively. In game 6, there was no chance of gaining, and some chance of losing, if you played uncooperatively, and in game 7, a player who failed to cooperate always received a smaller payoff, whatever the partner did. The partner's play affected only the amount that the player lost. Despite this, in every game but the last, uncooperative play predominated. Even in the last game, the outcome was very close to what you would expect if the players had picked their strategies by tossing a coin. Moreover, as the game progressed, the tendency to cooperate became weaker rather than stronger.

Why players fail to cooperate is not altogether clear. A player might want to exploit a partner or might fear that a partner is about to exploit him or her. A player might not understand what the game is all about or might doubt that the partner does—though this last possibility doesn't seem very likely. If a player with insight into the game fails to cooperate because she fears her partner won't "read her message," she would presumably cooperate when she was permitted to talk to her partner and make her message plain. But, in fact, there was only a slight increase in cooperation when the players were allowed to communicate.

Not only were the players slow to cooperate with each other; they appeared to be almost oblivious to what their partners did. They were not even suspicious when their partner's play was identical to their own. Suspicious or not, their playing uncooperatively 60 percent of the time for fifty-trial periods in the face of this behavior borders on the incredible. *There was roughly the same amount of uncooperative play when the experimenter duplicated the subject's play as there was when the experimenter always played uncooperatively.*

The subjects seemed to regard these games as purely competitive: beating one's partner was most important, and the

player's own payoff was only secondary. This tendency to compete, which has been observed by many experimenters and which increased as the game progressed, has been attributed to boredom and the smallness of the monetary payoffs. One wonders if this determination to beat one's partner would diminish if an appreciable amount of money were at stake.

Brian Forst and Judith Lucianovic (1977) also reported on a prisoner's dilemma "game," one played with real prisoners. They compared the fraction of codefendants who either pleaded, or were ultimately found, guilty with that of solo defendants and found the two were substantially the same. But despite the realistic trappings, they doubted that the situation really was a prisoner's dilemma since the accused usually managed to communicate even if both were in jail while awaiting, and during the trial, and there was a very real prospect of later retaliation by the betrayed partner.

A Bidding Experiment

In 1963 James Griesmer and Martin Shubik ran some experiments using as their subjects Princeton undergraduates. The basic experiment went like this:

Two players simultaneously picked a number between 1 and 10. The player who picked the higher number received nothing. The player who picked the lower number received that amount (in dollars or dimes, depending on the experiment) from the *experimenter.* If both players picked the same number, they tossed a coin and the winner of the toss took that number. If, for example, one player picked 5 and the other picked 7, the player who picked 5 received 5 units from the experimenter, and the other player received nothing. If

both players picked 5, a coin was tossed and the winner of the toss received 5 while his partner received nothing. After each trial, the players were informed what their partners had done, and the game was repeated. The players were separated throughout the series and could not consult.

This game is almost identical to the gasoline price war example discussed earlier. By the same induction argument used before, we can "deduce" the superiority of the most uncooperative strategy: a bid of 1. If both players adopt this strategy, the payoffs will be very small—on average 1/2. If one player is competitive and the other is not, the competitive player will generally win something, but not very much. The greatest rewards clearly go to the players who cooperate. If both players keep their bids high, the joint profit will be greatest, and though one of the players will necessarily get nothing on any particular trial, repetition of the game gives both players a chance to do well.

If both players in a pair realize that they will do best if they cooperate, they still have the problem of figuring out how to synchronize their bidding, since no direct communication is possible. The most direct way to cooperate, and the way that maximizes the joint expected profit, is to bid 10 all the time. This yields an average return of 5 to each player. The one flaw in this plan is that a player can still get nothing, since in the case of ties the payoff is determined by chance. At a small cost in expected value, the players can raise their assured profit considerably. On the first of each pair of trials, each player picks a 10. The player who won the toss on the first trial picks another 10 on the second trial, while his partner chooses a 9. Thus, on every pair of trials, one player always wins 10 and the other always wins 9. The average return to each player is 4 3/4, and the guaranteed return is 4 1/2.

As it happens, almost all the players approached the game competitively; they tried to outsmart one another rather than

jointly exploit the experimenter. Some players made cooperative overtures by bidding 9's and 10's alternately, but these were most often misinterpreted as attempts to lull the partner into a false sense of security. (This was later confirmed by the subjects in explicit statements and was obvious from their behavior during the play.) A few of the pairs did manage to get together, and when they did, it was invariably by alternating bids of 9 and 10, so that each won 9 on every other trial. This is not the most efficient way to cooperate, of course, but it is fairly good for subjects playing for the first time.

One of the things the experimenters wanted to study was the "end-effect"—uncooperative play on the last of a series of trials—but no evidence of it was observed during the experiment. When players managed to get together, they played cooperatively throughout the series of trials. If anything, there was more cooperation at the end than at the start. The experimenters tried to isolate end effect by telling some pairs the exact number of trials they would play and not telling others. Presumably the end effect would occur when the players knew which was the last trial, but in fact that there was no difference in play. Perhaps the payoffs were too small to motivate defection. Generally, two players who "found each other" were so pleased that they left well enough alone. (When the players were uncooperative from the first—and most were—the situation didn't arise.)

One variation of the game is worth describing in detail because it illustrates a point I made before in connection with utility theory. In one experiment, the basic game was played in sets of three trials, with one additional rule: if in the first two games a player received nothing, in the third game the player automatically received whatever he bid. This meant that a player could *always* get a payoff of 10 by bidding 10 on all three trials.

Players who were competitive in the earlier games usually played competitively in this game as well. Each of the players made low bids during the first two trials, and each won one of the first two games. The situation then was exactly the same as if they had been playing an ordinary game, since the new rule was not applicable. The players realized that if they played competitively on the last trial, as they had before, their profit, even if they won, would be considerably less than it would have been if they had played differently from the start. The realization that they could have done better apparently affected their utility functions, for they bid much more in the third game than they had earlier—in what essentially were identical situations.

An Almost Cooperative Game

There is another kind of cooperative, two-person game in which the basic problem is coordinating both players' strategies in accordance with their mutual interest. In games in which the players cannot communicate directly, Thomas Schelling (1958) suggested, the players must look for certain clues to help them anticipate what their partners will do. The clue could be a past outcome, for example, or it might be suggested by the symmetry of the payoff matrix. Richard Willis and Myron Joseph (1959) ran experiments to test Schelling's theory that prominence was a major determinant of bargaining behavior. In all, three different matrices were used. I call them 1, 2, and 3 (see figure 5.25).

The players were divided into two groups. Group *A* played game 1 repeatedly and then switched to game 2, which it also played repeatedly. Group *B* started with game 2 and ended with game 3. No communication between players was allowed during the play.

Figure 5.25

GAME 1

(10, 20)	(0, 0)
(0, 0)	(20, 10)

GAME 2

(10, 30)	(0, 0)	(0, 0)
(0, 0)	(20, 20)	(0, 0)
(0, 0)	(0, 0)	(30, 10)

GAME 3

(10, 40)	(0, 0)	(0, 0)	(0, 0)
(0, 0)	(20, 30)	(0, 0)	(0, 0)
(0, 0)	(0, 0)	(30, 20)	(0, 0)
(0, 0)	(0, 0)	(0, 0)	(40, 10)

Clearly, the players must solve the problem together. Unless they pick the same row and column, neither player will get anything. As a secondary goal, each player may try to reach the most favorable outcome of those that lie along the diagonal.

In game 1, Schelling offers us no clue as to what the proper play should be. In game 3, also, no clear-cut outcome suggests itself, but one of the two middle strategies would seem more likely than the extreme ones. It is only in game 2 that symmetry dictates a clear choice: the (20, 20) payoff corresponding to each player's second strategy.

What actually happened was quite surprising. When group *A* played game 1, there was a battle of wills in which each player fought for dominance. Each player played the strategy that would give him a payoff of 20 if the other yielded. After the initial battle, in which there were few agreements, one player finally gave way and that team settled

down to an equilibrium point. When they switched to game 2, their play was strongly influenced by what had happened before. Rather than go to the (20, 20) payoff, which symmetry would suggest, these players went most often to one of the asymmetric equilibrium points. Three-quarters of the time, the player who was dominant in game 1 remained dominant in game 2 too.

The actions of group *B* were even more unexpected. They started with game 2 and arrived at an agreeable equilibrium point much faster than group *A* had, which is not too surprising. But the agreement tended to be at one of the extremes—the first row and column or the third row and column—rather than on the middle row and column as suggested by symmetry and Schelling. As before, when the players in group *B* switched to game 3, the dominance that had been established in the earlier game prevailed.

Generally, repetitive play—picking the same strategy again and again—was the signal that was used to suggest an outcome. The repetitive play continued when it struck a responsive chord in the partner and even, often, when it didn't. The most equitable outcome in games 1 and 3 would be a synchronized, alternating scheme in which one partner is favored in one play and the other is favored in the next. This never happened, however, probably because such an arrangement is too intricate to arrange without direct communication.

Mother Nature as a Strategist

We observed earlier that game-theoretical models have a remarkable versatility. A model fashioned for one purpose often turns out to serve several others as well. A surprising

example of multipurpose modeling is the application of game-theoretical concepts to evolution and ecology.

Game theory is usually thought of as a tool used by a thinking person involved in a game with other thinking people. But in a fascinating article, John Maynard Smith (1978) described some very unusual applications of game theory in which organisms "choose" very sophisticated strategies that enable them to survive as a species. These strategies are not chosen deliberately by individual organisms but are formed collectively by the entire species. It is as though there is an "invisible hand"—similar to the one that is understood to work in economics—that weaves individual behavior into a pattern for the entire species. These individual behavior patterns and their interactions may be described by payoff matrices; by analyzing these payoff matrices you can determine whether, and in what form, the species will survive.

The *fitness* of an individual organism is its ability to survive and have offspring—the ultimate goal in the survival game. The fitness of a species, also, is its ability to survive. These kinds of fitness are quite different and what favors one may not favor the other.

A potential conflict between these two kinds of fitness raises a basic question about the evolutionary mechanism. If the individuals of a species possess an altruistic trait, this may make the species fitter but it may make the individual that possesses it less so. A bird giving a warning cry when a predator appears helps the species to survive but, by calling attention to itself, hastens its own doom. If the bird remained mute, it would live longer and produce more offspring.

So the paradox comes to this: it is important for the survival of the species that individual members take certain risks, but it is the individual, not the species, that mates and passes on the traits it possesses. The bird that gave the warning cry will be dead, and its altruistic genes will be lost, while the

selfish bird that stood mute will pass on its more prudent genes. Future generations of unwarned birds will pay the price for this temporary victory of opportunism, or so it appears at first.

But altruistic genes do survive and make the species fitter; and William D. Hamilton (1964) found the reason why. He observed that the critical factor in the survival of a gene isn't really the welfare of the organism that possesses it; rather it is the preservation of the gene itself and its replications. If a bird, taking a stroll with its family, loses its life because it gives a warning cry but saves ten of its offspring that would otherwise be killed, the loss of its own genes must be balanced against those of its offspring, but properly discounted.

The "discounting" process goes something like this: Since an offspring of the bird is as likely to inherit one parent's genes as the other, it will contain a particular gene of one of its parents only half the time. From the gene's point of view, the survival of its possessor is equivalent to the survival of two offspring. A sibling would also be worth half a gene, and more distant relatives would have even greater discounts. Hamilton added up the bits and pieces of the genes that were likely to be saved; if they outweighed the genes that were lost, he concluded the trait would survive.

In the game-theoretical survival model, then, the payoffs are not in terms of the individual's survival but in terms of the survival of the genes. A natural question to ask now is "Precisely which genes will survive?" And this is the question that John Maynard Smith attempts to answer.

Imagine a world in which all the members of a species have a particular trait. Subsequently there is a mutation—a small number of organisms are introduced with an alternative trait. What determines whether this new trait modification will die out or prosper? Smith has constructed a mathematical model that allows you to determine whether a trait is sta-

ble—that is, whether alternative traits will die out if they are introduced. He describes his model in the context of male animals competing for the attentions of a female.

There are many situations in which animals, like humans, have conflicting interests. When two animals compete for a mate or for territory, they generally begin by making threatening noises and taking an aggressive stance rather than by engaging in direct physical contact. If the confrontation appears to escalate into a physical struggle, they have a choice: they can back off, forgoing whatever spoils are involved but surviving for another day, or they can engage in serious physical combat. Smith calls animals that avoid combat *doves* and those that engage in combat, *hawks*. When a hawk engages a dove, it will gain whatever prize is at issue without even having to fight. But if a hawk encounters another hawk, one or the other is likely to be killed or seriously injured.

Now suppose that in a population composed solely of doves a mutation occurs and the population acquires a small number of hawks. Initially the mutants will thrive, because in confrontations between hawks and doves the doves will retreat and the hawks will mate; doves will only mate when there is no competition or when they are competing with another dove less persistent than themselves. But the early success of the hawks contains the seeds of their own downfall; as the hawks begin to proliferate, confrontation between hawks will become more common. To a dove, confrontation with a hawk is discouraging since the dove always leaves empty-handed; but to another hawk, such a confrontation can be a disaster resulting in death or serious injury. You might guess that at a certain point the doves and hawks will reach an equilibrium. This turns out to be the case. Smith assigned quantitative payoff values to the various possible confrontations, and on the basis of these payoffs, he tried to predict the long-run outcome.

Smith assumed that the animal that eventually mated received a payoff of +10 and an animal that was seriously injured received a penalty of −20. Two doves in combat do not hurt one another but since they do waste considerable time trying to bluff one another, Smith penalizes winner and loser alike by assigning them a payoff of −3. (Apparently when hawks engage either doves or other hawks, the matter is settled quickly.)

Putting these pieces together, Smith derives the payoff depicted in figure 5.26, which reflects the various possible confrontations.

Figure 5.26

	HAWK	DOVE
HAWK	(−5, −5)	(+10, 0)
DOVE	(0, +10)	(+2, +2)

When two hawks meet, one will gain 10 (+10) and the other will lose 20 (−20), so on average a hawk gets (−20 + 10)/2 = −5. When two doves meet, the successful dove gets +10 − 3 = 7 and the unsuccessful dove gets −3, so on average the payoff is 2. At a mixed confrontation the hawk gets 10 and the dove loses nothing.

In general when an organism of type X meets an organism of type Y the *expected payoff to* X is denoted by $E(X, Y)$; this payoff is the overall effect of wasted time, the risk of injury, and the chance of mating successfully and includes any other factor that is relevant to the spread of X's genes.

A gene I (or strategy I) is called *evolutionarily stable* if, whenever an alternative mutant, J, is introduced, I prevails and J dies out.

According to Smith's formal criterion, strategy I is evolutionarily stable if, for any alternative J, either condition 1 or condition 2 is true:

Condition 1: $E(I, I)$ is greater than $E(J, I)$
Condition 2: $E(I, I) = E(J, I)$ and $E(I, J) > E(J,J)$

Smith's reasoning goes like this. If a mutant, J, is introduced into a population in which I predominates, both I and J types will primarily be confronting I's. If an I does better against an I than a J does against an I, the I population will grow faster than the J population.This is expressed by condition 1. If both types do equally well against type I, both types will grow at the same rate and eventually there will be an appreciable number of type J organisms. If I's do better against J's than J's do against each other, this advantage will be critical and again, the growth rate of the I's will be greater than that of the J's. This is the effect of condition 2.

Let's see how this is applied to the conflict of hawks and doves. Suppose that a mutant dove is introduced into a population consisting only of hawks. Once doves gain a toehold, they will grow faster than the hawks because $E(D, H) = 0$ is greater than $E(H, H) = -5$. But this doesn't mean doves are fitter than hawks. If a mutant hawk is introduced into a population consisting only of doves, the hawks would grow faster than the doves since $E(D, D) = 2$ less than $E(H, D) = 10$. Neither a population of all hawks nor a population of doves is evolutionarily stable.

Smith observed that if the population adopts a "mixed strategy" so that eight-thirteenths of its population is hawks while the remaining five-thirteenths of its population is doves, it will remain evolutionarily stable against invasions from either hawks or doves. If we call such a population M, then it turns out that $E(H, M) = E(D, M) = E(M, M) = 10/13$; that is, everyone does the same against the mixed population initially. But if hawks proliferate, condition 2 takes effect, and since $E(M, H) = -40/13$ is greater than $E(H, H) = -5$, the mixed population will prevail. Similarly, if the

doves become too numerous, their numbers will be thinned since $E(M, D) = 90/13$ is greater than $E(D, D) = 2$. Instead of formulating this in terms of three different populations, you can also think of the hawks and doves as being in equilibrium; if any imbalance occurs, the original state will be restored.

It is not hard to verify that the mixed strategy just described is evolutionarily stable, but it is not clear how a population would implement it in practice. There are at least two possibilities, and either will serve. One possibility is that there is a hawkish gene possessed by eight-thirteenths of the population that prescribes that its possessors invariably act hawkish. The other five-thirteenths of the population would have a dovish gene, with an analogous effect on its possessors. On the other hand, it is conceivable that there be a single gene that every member of the population contains which prescribes hawkish behavior eight-thirteenths of the time and dovish behavior five-thirteenths of the time. Either population would be evolutionarily stable.

This analysis of evolutionary stable strategies was used to explain the behavior of male dung flies seeking a mate. It is well known—both to human biologists and male dung flies—that females lay their eggs in cowpats; male dung flies, hoping to mate, wait for them there. Since females prefer freshly made cowpats, the obvious strategy for the male is not to linger at any one cowpat but to stay only briefly and move on as new cowpats are formed. But if every male adopted the obvious strategy, competition would be intense: a male would do better to adopt the subtler strategy of remaining after the others have left and keeping the few female stragglers that happened along to itself. The fact is that *any* single pure strategy adopted by every male would be inefficient for the species as a whole and would lead to intense competition for some females and to the neglect of others. A much more

efficient strategy—and one that has been adopted in practice (see figure 5.27)—is to stagger the departure times so that the bulk of the males leave early in search of the preponderance of females but others leave more gradually and mate with the few tardy females. Although this mixed strategy has been observed in practice, it is still not clear how the strategy is implemented—whether each fly has its own fixed strategy and there is a mixture of different types of strategies or whether all the flies do the same thing: choose their own mixed strategy.

The ability of organisms to vary their behavior is often

Figure 5.27

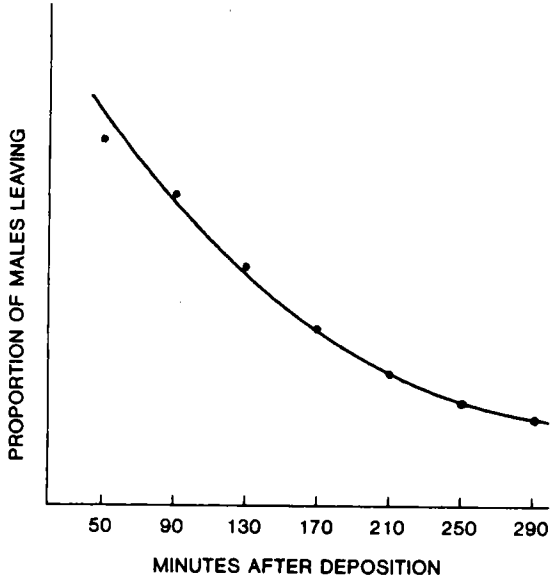

Relative Numbers of Male Dung Flies, *Scatophaga Stercoraria*, Leaving a Cowpat at Different Times After Its Deposition

NOTE: Reprinted, by permission of the publisher, from John Maynard Smith and G. A. Parker, "The Logic of Assymetric Contests," *Animal Behavior* 24 (1976):171.

beneficial to the species—this was true for male dung flies and it is also true in the hawk-dove model. One potential mutant in the hawk-dove model is particularly interesting; actually found in nature, it illustrates how nature resolves a paradox that plagues human beings, the prisoner's dilemma.

In the hawk-dove model, neither the hawkish nor the dovish strategy is completely satisfactory. The dove does well enough against other doves but does not get a fair share against hawks. Against doves, the hawk does as well as possible but confrontations with other hawks lead to disaster. What is needed is some sort of mediation that avoids physical combat yet prevents a conciliatory attitude from being exploited. One way of safeguarding what you have and still avoid excessive physical combat is to adopt what is termed the bourgeois strategy.

Suppose that in a population of animals, each has a certain region that he considers his own territory. From time to time there are conflicts (about a mate, perhaps) and once again we have the responses discussed earlier: after the preliminary posturing, the animal can flee (the doves) or can fight (the hawks). To these two old strategies we consider a new one: the bourgeois strategy. Bourgeois animals act like hawks when they are on their own territory but like doves when they are not. Since territory belongs to only one animal, in a bourgeois-bourgeois conflict there will never be any escalation; the stranger always backs off. Bourgeois-bourgeois duels are settled rapidly, with the intruder getting nothing and the inhabitant getting +10; a payoff of +5, on average. In the payoff shown in figure 5.28, it is assumed that each bourgeois has an even chance of being in his or her own territory.

Since $E(B,B) = 5$ is greater than both $E(H, B) = 2.5$ and $E(D,B) = 1$, the bourgeois strategy, once established, is safe against both hawkish and dovish incursions. Also, $E(B, H) = -2.5$ is greater than $E(H, H) = -5$, and $E(B, D) = 6$ is

Figure 5.28

	HAWK	DOVE	BOURGEOIS
HAWK	(−5, −5)	(10, 0)	(2½, −2½)
DOVE	(0, 10)	(2, 2)	(1, 6)
BOURGEOIS	(−2½, 2½)	(6, 1)	(5, 5)

greater than $E(D, D) = 2$, so either a hawkish or a dovish population would be vulnerable to a bourgeois invasion.

The predominance of the bourgeois strategy over hawks and doves seems to be confirmed in practice. Smith (1978) describes two examples—one observed by Hans Kummer and the other by N. B. Davies.

In a species of baboon in which a male forms permanent ties with one or more females, Kummer put male *A* together with a female while male *B* looked on. When the three were later reunited, *A*'s association with the female was not challenged by *B*. To avoid the possibility that *A* simply dominated *B*, several weeks later *B* was isolated with a female while *A* looked on. This time it was *A* who became reconciled to *B*'s association with the female. Possession was the cue for settlement of disputes.

Davies noticed that the speckled wood butterfly also adopts the bourgeois strategy in territorial disputes. Male butterflies seek sunlit spots on the forest floor because this is where the females are generally found. There are never enough spots for all the males so they are constantly flying in the forest canopy, scouting for vacancies. When a stranger invades an occupied spot, both stranger and resident circle briefly upward. One continues upward and the other—invariably the original resident—settles back on the spot. Davies noticed that possession is not merely a matter of strength; most of the butterflies he observed eventually settled in somewhere. Again, it was possession, not strength, that resolved these confrontations.

If a butterfly was unchallenged in a sunspot for more than a few seconds, it then considered the spot its own. Davies then asked, What if each of two males considered itself the owner? To answer the question, Davies insinuated an alien butterfly into a spot that already had an owner and waited. When the two butterflies confronted each other, they joined in a spiral flight that lasted ten times longer than the earlier ones in which it was clear to both parties who belonged and who didn't. A butterfly was clearly willing to escalate its quarrel on what it believed to be its own turf.

Not all territorial disputes are resolved as neatly as this; brute strength may play a role. When male fiddler crabs were observed disputing a burrow, the possessor was generally victorious (in 349 out of 403 cases, Smith reports). But in the minority of cases where the wandering male was successful, he was stronger.

A Computer Simulation of an Evolution Model

A few years ago Robert Axelrod, a political scientist, devised two interesting experiments involving people and computers (1980 *a,b*). At first glance the experiments seem to have little to do with Smith's model, but, in fact, the experiments and the model are very closely related. It turns out that the experiments yield insights into a number of different aspects of the evolutionary process.

In each of these experiments a number of people who were familiar with game theory—among them evolutionary scientists, social scientists, mathematicians, and so forth, many of whom had published papers in the field—were invited to enter a prisoner's dilemma tournament. The particular prisoner's dilemma involved had the payoffs shown in figure 5.29.

Figure 5.29

		OPPONENT: COOPERATE	OPPONENT: DEFECT
PLAYER	COOPERATE	(3, 3)	(0, 5)
	DEFECT	(5, 0)	(1, 1)

The participants played a round-robin tournament in which every entrant played every other entrant a series of about 200 prisoner's dilemma games with the payoffs shown in the figure. The players also played one round with a player who played randomly and another with the mirror-image of themselves.

Once the tournament started, the players no longer participated directly. The players had to convey their playing strategies to a computer in the form of a program that was created in advance and that would remain unchanged throughout the tournament. The program indicated what the player wished to be done at each turn. This decision might be made on the basis of how the player's partner had acted earlier, it might be to act randomly, or it might invariably be to defect or cooperate. The players knew in advance that they would confront a random player, but they had no idea what the other player would do. (They did know that the other players were sophisticated, however.)

The players had a wide range of possible strategies from which to choose. At one extreme, they might always defect. While this strategy may seem intuitively unappealing, the arguments that support it are hard to refute even when the game is played repeatedly, as it was here. A defection strategy does best against strategies that ignore past experience—other purely competitive strategies, purely cooperative strategies, and random strategies. At the other extreme is the pure cooperative strategy, which always turns the other cheek. This strategy does fine when it encounters others of its kind

but it doesn't do so well otherwise. All the other strategies fall between these two extremes.

Most of the strategies that were selected responded in kind to a cooperating partner—if not at first, eventually. The most important difference between the strategies was how they dealt with defecting partners. Some vindictive strategies defected and continued to defect until the end of the match if they observed a single defection by partner, whatever partner did after that defection. Some defected on the next play and then cooperated if their partner saw the light. And some continued to cooperate for a few plays, giving their partners a chance to reform before retaliating.

Some of the programs played a more dangerous game. They defected once or twice in the hope that they had a tolerant partner who would "educate" them before retaliating seriously; when they did this successfully, they picked up a few extra points. Some of the more cynical programs defected at the end of a run (in the first tournament, which was of fixed length), exploiting the fact that there are no rewards or punishments once the run is over. And finally, some strategies culled their partners' past history for cues to their future behavior. If their partners tolerated defections in the past, they would take a chance but not otherwise.

Considering the sophistication of the players and the variety of strategies available, the outcome of the experiments were surprisingly clearcut: "nice strategies"—those that never defect unless their partner defects first—invariably outperformed the others. There were no exceptions to this rule; each nice strategy outperformed every strategy that wasn't nice.

The nice strategies that did best punished defections but had short memories. They always gave their opponent a chance to reform. Strategies that irrevocably punished defectors for the rest of the run did not do nearly as well. When an

unforgiving (but nice) strategy ran into an occasional defector trying to pick up a few extra points, there was a disaster for both. When the defector's misdemeanor was treated as a felony he lost heavily; but reform by the defector was precluded by the unyielding strategy of the partner so that ultimately the nice strategy lost as well.

The winner of both tournaments was a surprisingly simple strategy submitted by Anatol Rapoport called Tit-for-Tat. The strategy cooperated on the first play; thereafter it did on each play what its partner had done on the previous play. Tit-for-Tat had a number of weaknesses: for one thing, it never "noticed" that it was playing a random player so it kept up a futile attempt to reform it. (This strategy did worse against the random player than any other program.) But despite this blind spot, and despite the existence of competitive programs that had the same general philosophy but sophisticated variations as well, the Tit-for-Tat strategy was the clear winner in both tournaments.

The three properties that made it a successful strategy were niceness, forgiveness, and provocability: it never was the first to defect, it didn't hold a grudge too long, but it did not ignore its partner's defections. Tit-for-Tat is similar to the bourgeois strategy in Smith's model: it cooperated (the bourgeois deferred to its partner when it was appropriate), but it did not let defections pass unpunished (the bourgeois escalated when its partner broke the territorial rules). Axelrod felt that choosing a survival strategy in Smith's model was much the same as choosing a strategy in the computer tournament. In fact, the computer tournament might be considered a model for ecological development. ("Ecological" rather than "evolutionary" because no mutations are allowed in the computer tournament—all the strategies are there from the start.)

Imagine a population of animals in which pairs confront one another from time to time. Assume that each animal has

a choice of two strategies—cooperate or defect—and that the payoffs are in terms of "expected number of surviving offspring" units. Imagine further a population in which each strategy is equally represented initially, and assume that the number of animals using each strategy in the next generation reflects the success that strategy had in this one (that is, the number of points it accumulated). Axelrod and Hamilton (1981) calculated what would happen if the strategies that were competing in the tournament met again and again, round after round. As the populations changed from generation to generation, a strategy's performance might change as well—what is fit in one population might not be in another. But in the main, it turns out, unfit strategies in the first round died out early and an equilibrium of sorts was achieved. By the five hundredth generation, eleven strategy groups had grown larger—the eleven strategies that had done best on the first round. Tit-for-Tat did best on the first round and continued to do best thereafter.

Assuming there is an analogy between surviving in a computer tournament and finding an effective evolutionary strategy, Tit-for-Tat's success in one environment suggests that it might be successful in the other as well. And it turns out that in the real world, it is, at least under some conditions.

Biologists have observed that many potential symbiotic relationships—relationships in which animals cooperate to their mutual advantage—exist in nature. But before these relationships can materialize, the same kinds of barriers to the cooperative solution discussed earlier, when players were people, must be overcome. Robert Axelrod and William D. Hamilton, an evolutionary biologist, discuss some of the barriers in their paper "The Evolution of Cooperation" (1981). Their statement—"The problem is that while an individual can benefit from mutual cooperation, each one can also do even better by exploiting the cooperative efforts of others. . . .

With two individuals destined never to meet again, the only strategy that can be called a solution to the game is to defect always despite the seemingly paradoxical outcome that both do worse than they could have had they cooperated" (p. 1391)—is a familiar story by now. When animals are unlikely to meet more than once, there is little motivation to cooperate, and Tit-for-Tat (or any other kind of cooperative strategy) is unlikely to thrive. When the prisoner's dilemma game is played repeatedly because the probability of meeting again is high, the Tit-for-Tat strategy comes into its own. And under these conditions, the strategy is often found in nature.

The variety of situations in which this kind of game-theoretical analysis can be applied is very surprising. For one thing, the organisms involved need not have a brain. As Axelrod and Hamilton point out, even bacteria have the capacity to play, because they can detect the response of a partner and then respond to their chemical environment. This is sufficient to allow them to punish defections and reward cooperation, because they have the ability to make their partners less fit, just as their partners have the same power over them. Moreover, their offspring can inherit this ability to discriminate.

The two organisms involved in a cooperative game don't even have to be able to "recognize" their partners; it is enough that they have continuous contact. Axelrod and Hamilton give a number of examples of such symbiotic relationships such as a hermit crab and its sea-anemone partner, a cicada and the microorganismic colonies of its bodies, and a tree and the fungi that inhabit it.

A more commonplace example of the Tit-for-Tat strategy in nature is the game played by the female fig wasp and a tree. The fig wasp sees its purpose as laying eggs; from the tree's point of view, the wasp's purpose is to pollinate flowers. If the fig wasp concentrates on laying eggs in a fig and neglects its pollinating duties, the tree cuts off the developing fig and all of its offspring die.

Axelrod and Hamilton cite another example of the cooperative game played by a small fish and its potential predator. The small fish eats parasites from the body, and even the mouth, of the large ones, while the large fish seeks its supper elsewhere. It is essential for such a relationship that there be consistent contact; this is managed by having an arranged meeting place. Such relationships have only been observed near the coast or on reefs but never in the open sea.

To maintain a cooperative relationship in a sequence of prisoner's dilemmas, there must be a good chance that defections will be punished and cooperation rewarded. In ant colonies that are stationary, symbiotic relationships are common; in honeybee colonies that are constantly moving, they are unknown. If there is a good prospect that the relationship will break up tomorrow, defection will become more likely today. Axelrod and Hamilton observe that seemingly harmless bacteria inhabiting the gut may suddenly become harmful when the gut becomes perforated. Other bacteria that are normally benign suddenly become dangerous when their host becomes sick or elderly. In each case the present reward outweighs the future prospects.

A final example is a game constructed by Martin Shubik as a casual amusement; it turned out to have several serious applications. Some time ago Shubik (1971) described the following parlor game: Auction off a dollar bill to an audience of several people but vary the usual rules. Allow the person who bids the most to pay what he or she bid and keep the dollar but require the person who made the *second highest bid* to pay what he or she bid without any compensation. Also, isolate the members of the bidding audience so that there is no collusion.

Shubik predicted that inexperienced players might bid

more for the dollar than it was worth. If your bid of 95¢ was raised to a dollar by a competitor, your choice would be to pay 95¢ without compensation (if the bidding dies at a dollar) or to bid $1.05 with the hope that your loss will be held to 5¢—the $1.05 you bid less the dollar you get. (The trouble is that your competitor may now bid on for the same reason.)

After Shubik made his predictions, Richard Tropper (1972) formally tested them. He had subjects participate in three different auctions (following Shubik's rules). In both the first and second auctions the "winning" bid exceeded the value of the prize. (In one case it was three times the size of the prize.) The subjects seemed to learn from experience, however, because in the last trial the bids were reduced substantially.

At first glance this game seems frivolous and without any application (not even to real auctions where only the winning bidder has to pay). In fact, however, it has several applications. When Shubik originally described this game, he mentioned that it was similar in many respects to an arms race. When two bidders in a Shubik auction have another round of bidding, both increase their potential losses but neither has increased his or her chance of winning the auction. In an arms race, similarly, the relative positions of both parties after they have increased their arms level may be the same yet they have both made a substantial extra investment.

The same model emerges in an evolutionary competition. John Haigh and Michael Rose (1980) use just such a model to describe a game involving two competing animals. Two adversaries interested in the same territory, or mate, resolve their quarrel by trying to outlast one another. They continue posturing until one loses patience; the more durable of the two obtains the reward but *both* pay the price in lost time.

After making certain quantitative assumptions about the payoffs, Haigh and Rose attempt to derive an evolutionarily

stable strategy. It is pretty clear that no single waiting time is optimal, for if there were a population that consisted only of animals with a fixed waiting time, a mutant that waited just a bit longer would prove superior. The authors conclude (what nature concluded before them) that a mixed strategy is called for. An even better solution would be to pick out some assymetric feature of the situation—superior size, longest in the territory, superior strength, and so forth—and use that to decide the issue; that would save everyone's time.

Some General Observations

The purpose of studying experimental games is to isolate the factors that determine how players behave. Ultimately, we hope, enough knowledge will be obtained to allow us to predict behavior. So far, however, there has been only a limited amount of experimentation, and the results have not been entirely consistent. Much of the work has centered on the two-person, non-zero-sum game, focusing on the prisoner's dilemma and in particular on identifying the elements that determine whether a person will compete or cooperate.

Participants in a prisoner's dilemma game, it has been established, often fail to cooperate. Several possible reasons for this were mentioned earlier. Players may be interested only in furthering their own interests. They may be trying to do better than their partner (rather than maximizing their own payoff). They may feel that their partner will misunderstand cooperative overtures, or, even if the partner understands, may only exploit them. Or they may simply not understand the implications of what they are doing.

There is little doubt that players often fail to see all there is to see in a game. Rapoport (1962) has suggested that people

fail to play minimax in zero-sum games (and they frequently do) because of lack of insight. There is some evidence that people fail to cooperate in non-zero-sum games for the same reason. Often players just do not conceive of playing other than competitively. In one set of prisoner's dilemma experiments, it was established by postexperimental interviews that, of twenty-nine subjects who understood the basic structure of the game, only two chose to play noncooperatively. In the bidding experiments conducted by Greismer and Shubik (1963 *a,b,c*) discussed earlier, many of the players were competitive because they weren't aware of all the possibilities. They perceived the game as being competitive and never thought about cooperating. And those who did cooperate felt uneasy—as if they were cheating the experimenter by conspiring illegally.

Despite all this, and even granting the truth of Rapoport's observations on zero-sum games, something more than the ignorance of the players seems to be involved when players fail to cooperate in non-zero-sum games. For one thing, highly sophisticated people often play uncooperatively, especially in one-trial games. For another, if the rewards are sufficiently enticing, either cooperative or competitive behavior can be induced. Apparently the factors that determine whether the play will be cooperative or competitive are quite complex and the experiments do not indicate the ratio of distrust, ambition, ignorance, and competitiveness that leads to competitive play.

Not all the significant variables in a game are controlled by the experimenter. One of the most important elements of the game is the personality of the players. It is to be expected that two people will react differently even in identical situations. If we are to have any chance of predicting the outcome of a game, therefore, we must look beyond the formal rules and into the attitudes of the players—a difficult job.

Experimenters have long been aware of the importance of personality and have tried to measure it or control its effect in a number of ways. The following clever scheme, for example, was devised by Jeremy Stone (1958) to measure the aspirations (or greediness) of players and their attitude toward risk.

Players were given a large number of cards and were told that they would be playing a series of games against an unnamed opponent. On each card were listed the rules of a particular game. A typical card might read: "You and your partner will each pick a number. If your number plus twice your partner's does not exceed 20, you will each get the number of dollars that you picked; otherwise you will each get nothing." Among the cards there would also be a game exactly like this one, but with the players' roles reversed. It would read: "You and your partner will each pick a number. If twice your number plus your partner's does not exceed 20, you will each get the number of dollars that you picked; otherwise you each get nothing." The cards were then paired, so that subjects played the role of both players—in effect, playing against themselves. The final score of a player was dependent on his or her own attitudes only.

Some experimenters tried, but failed, to correlate what players did in the prisoner's dilemma with their scores on certain psychological tests. Some connection *was* found between the players' political attitudes and how they played. In one study by Daniel R. Lutzker (1960), players were scored on an internationalism scale on the basis of how they responded to statements such as: "We should have a world government with the power to make laws which would be binding on all member nations," or "The United States should not trade with any Communist country." Lutzker observed a tendency of extreme isolationists to cooperate less in the prisoner's dilemma than did extreme internationalists.

Morton Deutsch (1960*b*), in a similar study, found that two characteristics related to the prisoner's dilemma were dependent on each other, and each of these, in turn, was dependent on the person's score on a particular psychological test. The traits Deutsch was concerned with are "trustingness" and "trustworthiness." To study these, a variation of the prisoner's dilemma was tried in which the play was broken down into two stages. In the first stage, one player picked a strategy. In the second stage, the other player picked a strategy, *after being informed what the first player had done.* The subjects played both roles, and in many different games. Players were *trusting* to the extent that they played cooperatively when they played first; they were *trustworthy* to the extent that they played cooperatively at the second turn in response to a partner's cooperative play. It turned out that, to a significant degree, "trustingness" and "trustworthiness" are positively correlated, and each trait is inversely related to authoritarianism (as measured by the F scale of the Minnesota Multiphasic Personality Inventory).

There have been attempts to relate other attributes of individuals, such as intelligence or sex, to the way they play, but with little or no success. The indications are that people's previous experience—and, in particular, their profession—affects their play. I have already mentioned how differently businessmen and students behaved in the last trial of a series of wholesaler-retailer games. Incidentally, a study was made of the prisoner's dilemma using prisoners as subjects. The results seem to contradict the adage about honor among thieves. Indeed, prisoners behaved very much like college students when playing this game: by and large, they tended not to cooperate.

Experimenters have also tried to control personality in games by imposing an attitude on the players. The players were told: "Do the best for yourself and don't worry about

your partner," or: "Beat your partner." In this way, it was hoped, the players' attitudes would be fixed. These instructions had little effect, however.

Behavior Patterns

Certain consistent patterns have been observed by several experimenters who have studied sequences of prisoner's dilemma games. There is a tendency for players to become less cooperative rather than more cooperative as the games are repeated—though it is not clear why. Also, the reaction of subjects to their partner's behavior is in a way consistent. Suppose a sequence of prisoner's dilemmas is divided into two halves. Suppose, also, that in one case a player's partner always acts cooperatively in the first half, and in another case the partner always acts uncooperatively in the first half. (In such experiments, the "partner" is either the experimenter or an assistant.) Uncooperative playing in the first half of the games, it was found, is more likely than cooperative playing to elicit a cooperative response in the second half.

The effect of allowing the players to communicate is also subtler than might at first appear. Some experimenters found that the chance of a cooperative outcome was increased if the players were allowed to communicate. An experiment by Deutsch (1960*a*) indicated, however, that this was so only for individualistic players: players who wanted to win as much as possible and who didn't care how their partners did. The cooperative player who was also concerned about how a partner did, and the competitive player whose main concern was doing better than a partner, played the same way, whether they were able to communicate or not.

In another experiment carried out by Deutsch (1960*a*) to

analyze further the effects of communication, three basic situations were set up: a bilateral threat, in which each player had a threat against the other; a unilateral threat, in which only one player had a threat; and a no threat. (People have a threat if they can lower their partner's payoff without changing their own.) In each of these three situations, some players were given the opportunity to communicate and others were not. Players who had the opportunity to communicate chose not to do so, however. In one situation—the bilateral threat—there was more cooperation when the players could communicate but didn't than when they did. When the players did try to communicate, negotiations often degenerated into repetition of threats. In the unilateral case, on the other hand, if the players communicated they generally became less competitive; and when the players could communicate but didn't, conflicts between them were created or aggravated. In short, the effect of allowing communication depends on the attitudes of the players, and, in turn, the attitudes of the players may be affected by the ability to communicate.

The most significant drawback to these experiments is, again, the insignificance of the payoffs. This was manifested explicitly and implicitly in many ways. It was reflected in the increased competition: beating one's partner became more important than maximizing one's payoff. It was indicated by the players' own statements; they admitted playing frivolously during long, monotonous, prisoner's dilemma runs. It was shown by the general tendency (at least on the part of students) to be cooperative in the wholesaler-retailer game once a tacit understanding had been reached and by the speed with which defectors were punished. (It is easier to punish a greedy partner when it costs you $10 than when it costs $10,000.) My point is not that the experimenters erred, but rather that the results must be analyzed cautiously.

The Two-Person, Non-Zero-Sum Game in Practice

It is difficult to study "games" that are actually "played" in real life. Still, it has been tried. Determining how fruitful these attempts have been is better left to the experts. I will be content to say just a few words in passing.

One of the fields in which game theory has been studied is business, and the analysis has been essentially descriptive. One study that I referred to earlier concerned the tactics used in a taxicab rate war. Elsewhere, competition in business was compared with armed conflict, and the conditions that tend to precipitate price wars were pinpointed. In the field of advertising, more formal models have been constructed, encompassing many problems: setting the advertising budget, allocating funds to the various media and/or geographic areas, determining the best time slots to advertise in, and so forth. There has even been an attempt to fit Indian burlap trade into the framework of game theory, allowing for possible coalitions, strategies, and threats.

Political scientists have also borrowed concepts from two-person, non-zero-sum game theory. I mentioned Maschler's application to disarmament models. Other models have centered about nuclear deterrence and bomb testing. Some models have been studied in extensive form, with the "plays" being military resources: transforming nonmilitary resources into military resources, standing pat, and so forth. The role played by communication, especially between the United States and Soviet Russia, has also been studied on this basis.

In these models, it is virtually impossible to get an accurate quantitative expression that will reflect the payoffs. Still, enough insight may be obtained by assigning approximate figures to justify the studies. In the next chapter, on games in which there are more than two players, even more difficult

quantitative estimates will be made—for example, the power of a member of a voting body will be examined.

Solutions to Problems

1. This game is the celebrated prisoner's dilemma, of course, and both the single-play and multiple-play cases were discussed in the text at length.There is really no clear answer to (a), (b), or (c), but Tit-for-Tat was the most successful strategy in the computer game.

2. a. You certainly can be better off if your strategies are restricted; a liquor retailer who lowers prices to meet the competition might well profit from a law that fixed prices.

Figure 5.30

		YOUR PARTNER: *A*	YOUR PARTNER: *B*
YOU	*C*	(1, 3)	(5, 1)
	D	(0, −90)	(3, −100)

b. If there is communication and the opportunity to form binding agreements, you can threaten to play *D* unless your partner agrees to play *B*; with no communication, your partner can play *A* and be reasonably certain you will play *C*.

c. In the symmetric game shown in figure 5.31, the outcome isn't clear. (It depends on whether you can communicate and give side payments and on your negotiating ability.) But if the rules of the game force you to move first and you play *C*, you will very likely get 30, your best possible payoff.

Figure 5.31

		YOUR PARTNER: *A*	YOUR PARTNER: *B*
YOU	*C*	(10, 10)	(30, 25)
	D	(25, 30)	(20, 20)

d. It can be to your advantage or disadvantage to have your opponent know your utility function. If goods are sold by bargaining and you want something badly, you'd best hide your feelings; but, as mentioned earlier, a company having financial difficulties would want to inform a labor union with which it is bargaining of its troubles.

3. The answer to all parts of this problem depend on your individual reactions; but there is no doubt that at some point, if your share is small enough, you will *not* try to maximize your profits but, instead, minimize your opponent's.

4. As the matrix stands, it is cheaper for the community to ignore the laws then to enforce them completely. But if the community enforces the law 10 percent of the time, the payoff matrix can be depicted in figure 5.32. With enough publicity, this payoff will induce rational speeders (?) to slow down at a cost of only 2 to the community.

Figure 5.32

		COMMUNITY: 10% ENFORCEMENT
DRIVER	SPEED	(−10, −7)
	DO NOT SPEED	(0, −2)

5. a. This is another form of the prisoner's dilemma. Notice that i dominates $i + 1$; that is, you do at least as well playing 9 as you do playing 10, at least as well playing 8 as you do playing 9, . . . It would seem, therefore, that each player should play 1. But it's hard to accept an outcome in which each player gets 1/2 on average when each player could have received 5. And it's harder yet if the game is repeated fifty times.

b. In an early experiment by James Griesmer and Martin Shubik (1963), players sometimes alternated 10's and 9's to signal their intent to cooperate. Their partners, more often than not, viewed this as a competitive game and felt their partners were either stupid or trying to deceive them; in a few cases the signal was read correctly to the benefit of both players.

6. I hope you haven't guessed it, but your partner was you. If you compare graphs *a, c; d, g; b, e;* and *f, h,* you will see that they are really the same game but from the point of view of different players. In effect, then, you were playing against yourself.

As a rough guide to how you did, I list the Nash value (obtained by maximizing XY) next to each of the paired games. If you received a lot of zeros, you're too greedy; if your scores were much lower (but positive) than the Nash sums, you're too modest.

GAME PAIR	SUM OF NASH VALUES
a-c	25 + 10
b-e	50 + 5
d-g	11.5 + 53.3
f-h	56.6 + 28.3
	239.7 for all 8 games.

This experiment was originally devised by Jeremy Stone (1958).

6

The n*-Person Game*

Introductory Questions

As we move from the earlier chapters to the later ones, players seem to lose control over their fate. In two-person, zero-sum games with equilibrium points, players could always exact their due and if there were no equilibrium points they could obtain it on average. In two-person, non-zero-sum games, they had to share control of their own fate with a partner but they had control over their partner's fate, which they could use as a threat. In *n*-person games, even this threat is generally denied the players. They must form coalitions with others and consider what inducements they must offer and accept. Once again, start by trying to answer these introductory problems.

1. An engineer, a lawyer, and a treasurer lose their jobs when their firm goes bankrupt. They are offered new jobs and an additional bonus if they join as a team. The additional bonus (in thousands of dollars) for the engineer-lawyer pair is 16, for the engineer-treasurer pair, 20, and for the lawyer-treasurer pair, 24. If all three join as a team, the bonus is 28.